

AUUGN

The Journal of AUUG Inc.
Volume 20 • Number 3
August 1999

Features:

Cluster Computing Components, Tools and Systems	11
My Home Network	25

Conference News:

AOSS-1: How it all went	7
AUUGwet: A post-mortem	8
Update: Open Source AUUG'99	29
What's on at USENIX	29

Regulars:

Book Reviews	27
From the Pages of Unigram X	35
Unix Traps and Tricks	45



**AUUG Membership and
General Correspondence**

The AUUG Secretary
PO Box 366
Kensington NSW 2033
Telephone: 02 9361 5994
or 1800 625 655 (Toll-Free)
Facsimile: 02 9332 4066
Email: auug@auug.org.au

AUUG Management Committee

Email: auugctee@auug.org.au

President:

David Purdue
David.Purdue@auug.org.au
Tattersall's
787 Dandenong Road
East Malvern VIC 3145

Vice-President:

Mark White
Mark.White@auug.org.au
Compaq Computer Asia/Pacific Pty Ltd
Level 8, Central Plaza II
66 Eagle Street
Brisbane QLD 4000

Secretary:

Stephen Boucher
Stephen.Boucher@auug.org.au
The Australian Industry Group
380 St. Kilda Road
Melbourne VIC 3004

Treasurer:

Luigi Cantoni
Luigi.Cantoni@auug.org.au
STM
Suite 3, 77 Mill Point Road
South Perth WA 6153

Committee Members:

Malcolm Caldwell
Malcolm.Caldwell@auug.org.au
Northern Territory University
Casuarina Campus
Darwin NT 0909

Lucy Chubb
Lucy.Chubb@auug.org.au
Aurema Pty. Ltd.
79 Myrtle Street
Chippendale NSW 2008

Günther Feureisen
Gunther.Feureisen@auug.org.au
GFH Computer Services
PO Box 1350
Double Bay NSW 1360

Peter Gray
Peter.Gray@auug.org.au
Information Technology Services
University of Wollongong
Wollongong NSW 2522

David Newall
David.Newall@auug.org.au
Tellurian Pty Ltd.
272 Prospect Road
Prospect SA 5082

Michael Paddon
Michael.Paddon@auug.org.au
Australian Business Access
33 Lincoln Square South
Carlton VIC 3053

AUUG Business Manager

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

Editorial

Günther Feureisen
Gunther.Feureisen@auug.org.au

It's conference time; not long to go now! Don't forget to get your registrations in!!

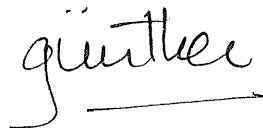
One of the biggest advances in high-end computing and super computing is the ongoing work into Clustering of Computers. Our feature this issue is an article by Mark Baker and Rajkumar Buyya entitled "Cluster Computing Components, Tools, and Systems".

Frank Crawford brings us up-to-date with "his" Home Network; and some feedback from some of our members on "their" Home Networks.

Elections have come and gone, with a new look committee and a new General Committee Member: David Newall, a long time AUUG member and supporter. I'd also like to pass on thanks to Lucy Chubb, our outgoing President; although short, her term as AUUG President did see us through some very difficult times.


One of the changes at AUUG is how we keep track of our members; the transition has caused a few hiccups, which we are currently trying to rectify. For those of you who haven't received their membership cards, you should get them soon. Liz Carroll has been working very hard to sort it all out. If you do have any questions about your membership, please give Liz a call or drop her some email.

That's it for this time – enjoy the issue, and have a great time at the Conference!



❖

*Thanks to our
Sponsors:*



TELLURIAN

Contribution Deadlines for AUUGN in 1999/2000

Volume 20 • Number 4 – November 1999:	October 17th, 1999
Volume 21 • Number 1 – February 2000:	January 17th, 2000
Volume 21 • Number 2 – May 2000:	April 17th, 2000
Volume 21 • Number 3 – August 2000:	July 17th, 2000

AUUGN Editorial Committee

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

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

Editor:
Günther Feureisen

Sub-Editors:
Mark Neely
Jerry Vochtelloo

Public Relations and Marketing:
Elizabeth Carroll

AUUGN Submission Guidelines

Submission guidelines for AUUGN contributions can be obtained from the AUUG World Wide Web site at:
www.auug.org.au

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

AUUGN Back Issues

A variety of back issues of AUUGN are still available. For price and availability please contact the AUUG Secretariat, or write to:
AUUG Inc.
Back Issues Department
PO Box 366
Kensington NSW 2033

Conference Proceedings

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

Mailing Lists

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

Disclaimer

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

Copyright Information

Copyright © 1999 AUUG Inc.

All rights reserved.

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

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

President's Column

David Purdue
David.Purdue@auug.org.au

society n. - 1. organized and interdependent community.

(The Pocket Oxford Dictionary)

As the incoming President, I think my first column should be about AUUG, what I think it is and where I think it is going.

There is not doubt that AUUG has been through a difficult patch. The society had become somewhat commercialised, and while this worked very well in the 80's and early to mid 90's, allowing AUUG to be very successful, in the late 90's it saw AUUG become less relevant to its members, which precipitated a fall in membership.

And so the management committee examined what it is AUUG should be, and started working towards refocussing on what originally brought the membership together. This is reflected in the choice of theme for the 1997 conference: "Technical Solutions."

We have also looked at the finances of the organisation and restructured so as to ensure we spend more on member benefits and less on the actual running of the organisation.

We are still in a mode of change. As it was expressed at a recent management committee meeting, we seem to change the flavour of AUUG at every meeting. I do not think this is a bad thing.

In the 90's we have been faced with the social disease called Economic Rationalism, and I believe it will be a few years before we are able to assess the damage to society caused by this. Economic rationalism has caused almost everyone we work with to take a purely inward focus in their dealings with others. It seems that the first question that is asked when any joint activity is proposed is, "What's in it for me?"

I want the AUUG membership to remember that AUUG is a society, an interdependent community. I think the greatest member benefit provided by AUUG is the other members of AUUG. All of us have different experience and different skills, and I hope we can learn to share those skills and experiences with other members.

But what is AUUG about? Our current tag line, "UNIX and Open Systems Users," nearly but does not quite sum it up. Certainly, we are interested in "openness", however we choose to define it, and I expect to take on an advocacy role in that arena.

But I think more importantly AUUG members are people who work with technology, and work with technology because they have a passion for it. Expect me to expand on this theme in subsequent columns.

In the mean time I am somewhat focussed on our conference, to be held in Melbourne on the 6th - 10th September. We have quite a respectable program (I have to say that, I am the program chair), and we are examining Open Source from all angles. I hope that many of you can attend, and ask that even if you can not, please let any one who would be interested know about the conference. Look at:

<http://www.auug.org.au/winter/auug99/>

It just remains for me to extend a huge thank you to Lucy Chubb, the retiring President. Lucy has lead AUUG through one of its most difficult periods, and for most of the year was too caught up with the minutia of running AUUG to have a chance to make much progress with the big picture. Lucy hands over to me a much healthier organisation, and I thank her sincerely for the commanding effort she put in to get AUUG to this state. Lucy will carry on serving on the management committee as Immediate Past President.

See you at the conference!




Winner of the 1999 AUUG Membership Survey Competition

And the winner is: AdamJ@netspace.com.au

Adam has won a free registration to Open Source - AUUG'99 being held at the Carlton Crest Hotel, Melbourne, on September 8th to 10th.

Congratulations!





Open Source™

LINUX

PERL

TCL/TK

PYTHON

APACHE





DNS & BIND

SENDMAIL

Many of the most important innovations in the computer industry don't come from large companies. They come from independent developers freely sharing their program source code and collaborating globally over the Internet. This movement has come to be called Open Source.

When you want to learn what's happening in the latest technologies - Ask Someone Who Knows...

O'REILLY®



Distributed in Australia and New Zealand by



For more information on Club O'Reilly membership please call customer service at Woodslane on 1800 803 443

<http://www.oreilly.com.au>

Returning Officers Report

Results of the 1999 AUUG Election

Chris Maltby
chris@aurema.com

The result of the 1999 AUUG election is as follows:

Votes received:

88 formal
2 unsigned

President:

David Purdue 58 (elected)
Malcolm Caldwell 30

Vice President:

Mark White 50 (elected)
Malcolm Caldwell 33
-exhausted 5

Secretary:

Stephen Boucher 52 (elected)
Malcolm Caldwell 36

Treasurer:

Luigi Cantoni (elected unopposed)

Committee:

Malcolm Caldwell (elected unopposed)
Gunther Feuereisen (elected unopposed)
Peter Gray (elected unopposed)
David Newell (elected unopposed)
Michael Paddon (elected unopposed)

Returning Officer:

Chris Maltby 61 (elected)
Peter Chubb 24
-exhausted/informal 3

Assistant Returning Officer:

Peter Chubb (elected unopposed)

❖

Results of the 1999 AUUG Membership Survey

Question01:

How long (years) have you or your organisation been a member of AUUG?

not a member = 5 (8.06 %)
less than 1 = 5 (8.06 %)
1 to 2 = 3 (4.84 %)
3 to 4 = 18 (29.03 %)
more than 5 = 31 (50.00 %)
TOTAL responses = 62

Question02:

What is your age group?

under 25 = 3 (4.84 %)
25 to 34 = 24 (38.71 %)
35 to 44 = 31 (50.00 %)
45 to 54 = 3 (4.84 %)
55 or over = 1 (1.61 %)
TOTAL responses = 62

Question03:

How many years IT experience do you have?

1-5 = 6 (9.68 %)
6-10 = 10 (16.13 %)
over 10 = 46 (74.19 %)
TOTAL responses = 62

Question04:

What other associations are you a member of?

ACM = 12 (14.46 %)
ACS = 9 (10.84 %)
Decus = 3 (3.61 %)
IEEE = 12 (14.46 %)
ISOC = 6 (7.23 %)
Linux User Group = 12 (14.46 %)
PC User Group = 3 (3.61 %)
SAGE-AU = 15 (18.07 %)
USENIX = 2 (2.41 %)
Other = 9 (10.84 %)
TOTAL responses = 83

Question05:

What is your primary job function?

Consultant = 9 (14.52 %)
Engineer = 3 (4.84 %)
Executive/Senior Manager = 6 (9.68 %)
MIS Manager = 1 (1.61 %)
Professor/Instructor = 2 (3.23 %)
Programmer = 4 (6.45 %)
Systems/Network Administrator = 27 (43.55 %)
Technical Staff Member = 4 (6.45 %)
Other = 6 (9.68 %)
TOTAL responses = 62

Question06:

When purchasing hardware or software, would you be:

Key Decision Maker = 19 (31.15 %)
Recommend to Decision Maker = 35 (57.38 %)
Non-Decision Maker = 5 (8.20 %)
Other = 2 (3.28 %)
TOTAL responses = 61

Question07:

What would be the approximate annual expenditure of your company or section on Information Technology?

under 50k = 12 (20.34 %)
50k-100k = 9 (15.25 %)
100k-250k = 9 (15.25 %)
250k-500k = 4 (6.78 %)
500k-1M = 4 (6.78 %)
over 1M = 21 (35.59 %)
TOTAL responses = 59

Question08:

What is the primary business of your organisation?

Agriculture, Mining, Construction = 1 (1.64 %)
 Consulting = 6 (9.84 %)
 Education = 7 (11.48 %)
 Finance = 2 (3.28 %)
 Government and/or Military = 9 (14.75 %)
 Hardware Vendor = 3 (4.92 %)
 Manufacturing (non computer) = 1 (1.64 %)
 Public Service = 2 (3.28 %)
 Retail, Wholesale, Distribution = 1 (1.64 %)
 Software Development = 12 (19.67 %)
 Systems Integrator, VAR, OEM = 5 (8.20 %)
 Telecommunications = 2 (3.28 %)
 University or research = 7 (11.48 %)
 Other = 3 (4.92 %)
 TOTAL responses = 61

Question09:

How many employees at your company?

less than 10 = 11 (18.33 %)
 10-25 = 5 (8.33 %)
 25-50 = 9 (15.00 %)
 50-100 = 3 (5.00 %)
 100-300 = 1 (1.67 %)
 300-500 = 3 (5.00 %)
 500-1000 = 4 (6.67 %)
 over 1000 = 24 (40.00 %)
 TOTAL responses = 60

Question10:

What non-UNIX operating systems do you use?

None = 1 (0.52 %)
 Macintosh = 16 (8.25 %)
 Mainframe = 11 (5.67 %)
 MS-DOS = 22 (11.34 %)
 Netware = 17 (8.76 %)
 OS/2 = 3 (1.55 %)
 Windows 3.x = 16 (8.25 %)
 Windows 95/98 = 54 (27.84 %)
 Windows NT = 46 (23.71 %)
 Other = 8 (4.12 %)
 TOTAL responses = 194

Question11:

Please list the publications you read regularly.

AUUGN = 51 (36.17 %)
 IT Magazine = 37 (26.24 %)
 IT Newspaper = 44 (31.21 %)
 Other = 9 (6.38 %)
 TOTAL responses = 141

Question12:

Do you feel AUUG meets your needs?

Yes = 50 (83.33 %)
 No = 10 (16.67 %)
 TOTAL responses = 60

A letter from the Secretary

Dear AUUG Member,

NOTICE OF ANNUAL GENERAL MEETING

The 1999 AUUG Incorporated AGM will be held in conjunction with the AUUG'99.

The AGM will be held at:

*AUUG'99 Conference
 September 8-10, 1999
 Carlton Crest Hotel
 Queens Road
 Melbourne VIC Australia*

The meeting will start at 5:00pm on Wednesday, 8th September 1999. All AUUG members are invited to attend. Although not required to gain entry, it would be appreciated if you could bring your membership card or membership receipt.

AGENDA

1. Apologies
2. Minutes of the previous meeting
3. Returning Officer's report
4. President's report
5. Secretary's report
6. Treasurer's report
7. Ratify appointments of Management Committee members
8. Ratify appointment of Assistant Returning Officer
9. Chapter Technical Conferences for during 2000.
10. AUUG'2000
11. Any other business.

Sincerely,

Stephen Boucher
 Secretary
 AUUG Incorporated

❖

The Australian Open Source Symposium (AOSS-1):

How it all went ..

Con Zymaris
conz@cyber.com.au

INTRODUCTION

What follows is merely one individual's recollection of this one day event, held recently in Melbourne, Australia. This shouldn't be read as being the official coverage, as it is quite likely that I've missed important details, or perhaps skewed the presented talks in some way. I imagine I'll be sent corrections from the speakers or other attendees in due course :-)

ORGANISERS

The Australian Open Source Symposium (AOSS) was organised and supported by AUUG Inc, ISOC-AU, Linux Users Victoria and SAGE-AU. Primary sponsor was SCO.

More details about the event are outlined here:

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

GNU DEBUGGER PROJECT

The AOSS was held on Wednesday, 16th June, at the Australian Industry Group training rooms, on St. Kilda Rd. Melbourne. This turned out to be an excellent venue for the attendees, if a little complex to find :)

After some initial hiccups with the video-projection system, the order of talks was flipped to allow Andrew Cagney from Cygnus Solutions to give his presentation on the GNU Debugger project. While Cygnus is based in California, Andrew works from Melbourne. To me, this clearly demonstrated what I believe is an interesting side-benefit of open source software projects. Mainly, that there is substantial scope for organisations to procure their technical skills across international borders, increasing the range and size of their resource pool, and of the staff that they can attract. I assume this is so because many of the tools and development methodologies which have come to fruition via open source projects, such as CVS, web-based and distributed help-desk systems, command line tools, also lend themselves greatly to distributed software development in general. Added to this is the fact that the Unix-like operating systems are fully functional to remote users, which makes for easier integration testing by developers when code-merging occurs, as will be the case when geographically distributed developers are used.

Andrew gave an amusing presentation which

included discussion on the complexities involved in development and maintenance of the GDB, mention of vast array of CPU architectures and target platforms supported, the use of similar tools in the future development of Sony PlayStation II games, and the heuristics used to know when to deprecate support for lesser used platforms/CPUs. He also presented graphically the code increase of the GDB system, noting the large leaps in size at various milestones, such as when HP merged a substantial amount of code into the main source tree in the near past. During some Q &A, Andrew outlined that various techniques of modularisation of code for the various target platforms was being looked as a means to keep the overall size of the GDB down.

MY TALK

After Andrew, I gave talk titled The Diary of an Open Source Project. More details on this are available:

<http://www.cyber.com.au/users/conz/aoss.html>

LINUX SUPERCOMPUTERS

Next came a solid presentation from Rajkumar Buyya from Monash University, on the topic of Parallel Processing with Linux. This was a substantial overview of the current state of play with Linux clustering technologies. Various models for building high performance computers were presented, including Beowulf clusters. Numerous high speed, TCP/IP and non TCP/IP networking technologies were looked at, along with analysis of the advantages and disadvantages. Management of Linux clusters was discussed, including outlines of tools which made the cluster as a whole look like a single computing device. Advantages given for why Linux clusters are good vehicle for various high-performance computing needs included: reduction in acquisition cost, reduction in downtime (if one unit crashes, the remainder keep running,) reduced on-going maintenance costs, etc.

EDDIE: WEB SERVER LOAD BALANCING

Geoff Wong, a member of the Eddieware core team was up next. He talked about both the technical aspects of Eddie, as well as the political effort involved in getting Eddie's primary sponsor, Ericsson to release the code under an open source licence. Eddie is 'aimed at delivering a commercial grade, quality of service driven web server solution' Eddie has been written in the functional programming language, Erlang (www.erlang.org) In detail, Eddie functions by having specified servers operate as controllers, distributing incoming traffic across designated back-end servers. These controllers also keep tabs of the availability (and load) of the backend web server. Eddie currently comprises an

enhanced DNS server which provides load balancing and monitoring of site accessibility and an intelligent HTTP Gateway which provides site based load balancing, reliability, scalability, and quality-of-service.



Around about now, we had a break for lunch, and dined on some delicious gourmet rolls and a fruit/cheese platter. This gave a good chance to catch up to people and find out how other open source projects around town were going.

COLOUR IN AN OPEN WORLD

After lunch came an interesting presentation from Graeme Gill, on the technicalities of the output of Colour from applications in an open, portable, interoperable manner. Graeme focussed on introducing the primarily programmer-oriented audience to the physics and biology of colour and colour quantisation. He discussed the various schemes presently used for providing accurate metrics of colour on monitor screens and in printers for generating hardcopy. Graeme talked about the problems of closed, proprietary mechanisms and code for providing these metrics, and his wish to produce an open system to foster the wider spread of quality colour output systems. He described the colour quantisation library he has developed, which he is releasing under the LGPL, with the hope that it will be widely accepted and utilised, and thus assisting adoption of the open standard used. During the Q & A, someone raised the issue that this kind of library could be used to give The Gimp the kind of professional colour management functions it presently lacks.

BUILDING RPMS

The following talk was given by my colleague Richard Keech. He covered the nitti-gritty of Red Hat's Package Management System. Included in the discussion was why RPMs are useful, how they worked, how to build source, binary and platform neutral RPMs, and RPMs for non-Linux OSs.

SCO

Peter Laytham from SCO was up next. He wanted to give the audience SCO's perspective on Open Source. He discussed the advantages to SCO of having an ever-growing base of Unix skills and applications. He suggested the advantages were a 2-way street, in that that success of open source was made possible by the major reduction in hardware costs which have arisen through mass-market pervasiveness generated by closed-source software. Peter hinted that one area where open source software could use increased advocacy, was in calming corporate fears of stray code from non-open

source applications making its into way widely used open source applications, exposing the corporate users to some form of legal liability. During the Q & A session, this topic was examined, and only instances of GPL'd code finding their way into closed-source applications were noted, not the other way around.

PANEL SESSION: OPEN SOURCE LICENCES

After afternoon coffee & tea, there was a panel discussion on open source licensing. Included on the panel were Andrew Cagney, Enno Davids (from Metva and current president of AUUG-Vic) and I. Numerous (thorny) questions had been garnered from the audience, including various problems with the GNU library licences as they pertained to embedded systems; incorporation of code into a project which had no clearly defined 'owner', and which could cause the open source developer to fall legally foul if shown to be acquired from non-open source means; discussion on which licences the various developers chose for their products, and why. There was also discussions of how various licences would hold up under severe legal scrutiny or indeed in court. It was generally agreed that the GNU and BSD licences would have the best shot of coming through unscathed in any legal proceedings. There was also some discussion of the proposed U.S UCC - article 2B, renamed recently to UCITA (Uniform Computer Information Transactions Act.) This will allow closed-source vendors to do things like disable end-user applications remotely, and make reverse-engineering illegal.

PROJECT PITCHES

Finally, the floor was opened to members of the audience to come up onstage and pitch their open source projects and ask for potential members



Thanks go to Michael Paddon, Adrian Close and Susie Close, and numerous other AUUG and LUV folks for helping to organise this successful event.



AUUGWet99:

A post-mortem

Malcolm Caldwell
Malcolm.Caldwell@auug.org.au

Well, it has been and is gone - the annual conference of the NT chapter of AUUG. A quick mini summary is "it went according to plans". Problem is there were no real plans worth mentioning. That's not completely fair - there

were plans and we could have held a good conference. Unfortunately not enough real work was done.

In the past the NT chapter has run conferences that run the spectrum from successfully to not-so-successful. This one went worse than the others.

I think the NT chapter is suffering what AUUG in general is suffering. A few years ago Open Systems was the buzzword of the industry. Everyone wanted a piece of the action, and computing practitioners (I nearly used the word professionals) wanted to learn more and participate in events that promoted Open Systems. More importantly their managers were vastly interested in Open Systems and did not hesitate to send their staff to conferences focusing on such noble subjects. Almost as importantly, vendors also wanted to be seen to stand behind Open Systems.

Then the world moved. Suddenly people were not interested in Open Systems, and were instead involved in other activities (like working out why their exchange server keeps crashing).

The next big thing was the WWW. Everyone wanted a piece of the WWW, and many knew that Unix was the power behind the WWW. So people came to our conferences and convinced managers to come to our conferences etc. The WWW is still a big thing - but I think it is a big other thing to what AUUG is about.

At the same time as this happened there was a move within the industry to outsourcing IT staff. This is something we noticed particularly strongly here in the NT. In the past technical people convinced their management that they should be sent to a conference that went for a one/two/three days. Now they had to convince themselves, and their clients (and their wives?) , that it was more important to go to a conference then to spend those three days making bucks.

That all being said, AUUG is still very relevant and is a great venue for keeping up to date in a constantly moving industry.

All that has got next to noting to do with what went wrong with AUUGWet99. It is worth saying that the number of people attending the NT chapter conference has diminished. This is not why this conference went so badly.

What it comes down to is:

- We lost a hard working secretary
- We lost easy contact with some people the committee
- Members on the committee did not have enough free time to run a conference.
- There was no free time each week for the committee to get together

We had scheduled AUUGWet to be held at about our normal time. (In April). We got a month out and realized that as we did not have a pamphlet

we should postpone the conference. We postpone until late May. We got close to the new date, and sill no pamphlet had gone out. So we postpone again.

At this stage we had two interstate speakers who had tickets booked to come to visit us. We had paid for those tickets. We rescheduled the conference for 15th and 16th of July. Luckily our interstate speakers were able to change their tickets and were willing to come for the new dates. They both indicated that they would not be able to come any later than that.

In the end, the pamphlet never got printed or sent out. We emailed all the members and every list of IT people in the NT that we could. I personally was determined to run the conference at this new date, and recoup any losses we had, and make sure the NT chapter members got as much benefit from the expenses incurred so far.

We got to the week of the conference. At this stage, two days before the conference there had been no registrations. Many people told me that they knew of people who were going to register but there were no actual registrations. I was in a bind. Part of me wanted to cancel the conference and never have anything to do with the NT chapter again. Part of me wanted to hope for the best and just run with and see how things went.

The day before the conference there was 1 registration. The next problem was that many of the local speakers were not able to speak or had pulled out. We decided late on the day before the conference to run the conference in a different way. This is what happened. We still held the workshop we had planed and charged \$100 for it. We canceled all but our two interstate speakers for the second day and we charged nothing to attend those papers. We emailed everyone we could think of to tell them to come to the two free papers. We also asked the local speakers we did have if they would mind presenting at some of our monthly meetings.

6 people came to the workshop and about 10 came to the two papers.

The main thing we managed to do at the conference was hold an AGM. At the AGM we called for some new volunteers. I resigned as the Chairman of the NT chapter, and I was replaced by Mark Bradbury. We then filled the position of secretary. I ended up being voted in as the new treasurer. We also got a few volunteers to help on the chapter executive.

I am sure it is clear the conference was not a 100% success. It was probably not even a 90% success. It was also not a 100% failure.

We (and particularly I) have learned some important lessons for those contemplating running a chapter conference.

- All members of the organizing committee should have time free in their schedules to meet regularly in the months before the conference
- All members of the organizing committee must be reachable via email. And not via accounts on ISPs that they consult infrequently.
- There must be enough people on the committee to do the work
- Have a timeline with deadlines so you can pull out if things are going bad

Other NT specific lessons

- Hold the conference in the wet season so air tickets are of peak.

The NT chapter is now in a better position than it was a few months ago. We have people identified to speak at the next few monthly meetings. More importantly we have new people available to organize the next conference.



If you have any experiences using Linux that you would like to share with other AUUGN readers, drop us a line at:

auugn@auug.org.au

We'd love to hear from you!



Cybersource

ACN: 053 904 082

Cybersource has been a Professional Services consultancy, specialising in the areas of Unix, Windows and TCP/IP since 1991. Cybersource also offers accredited, professional-grade support for Red Hat Linux and other open source (free) software. Therefore, the last 'valid' reason for not taking advantage of great software like Perl, Linux, SAMBA and Apache has just disappeared. Organisations can benefit from the robustness, flexibility and value of open source software, and know they have an experienced team of IT professionals available to provide commercial-level support, when needed.

Contact us for full details.

Telephone: 03 9642 5997
URL: <http://www.cyber.com.au/>
Email: info@cyber.com.au

Cluster Computing Components, Tools, and Systems*

Mark Baker
mab@sis.port.ac.uk
Division of Computer Science
University of Portsmouth

and

Rajkumar Buyya
rajkumar@dgs.monash.edu.au
School of Computing Science & Software Eng.
Monash University

ABSTRACT

The availability of high-speed networks and increasingly powerful commodity microprocessors is making the usage of clusters, or networks, of computers an appealing vehicle for cost effective parallel computing. Clusters, built using commodity-of-the-shelf (COTS) hardware components as well as free, or commonly used, software, are playing a major role in redefining the concept of supercomputing. In this paper, we discuss the motivation for moving to COTS-based clusters for supercomputing applications. We describe the enabling technologies and present three state-of-the-art cluster-based projects. Finally, we summarise our findings and draw a number of conclusions relating to the usefulness and likely future of cluster computing.

1 INTRODUCTION

The recent advances in high speed networks and improved microprocessor performance is making clusters or network of workstations an appealing vehicle for cost effective parallel computing. Clusters built using commodity hardware and software components are playing a major role in redefining the concept of supercomputing leading to the use of clusters as commodity supercomputers. The use of workstation clusters⁺ to prototype, debug and run parallel applications is becoming an increasingly popular alternative to using specialised, typically expensive, parallel computing platforms such as the Cray T3E or the IBM SP2. An important factor that has made the usage of workstations a practical proposition is the standardisation of many of the tools and utilities used by parallel

* The extended and enlarged version of this article appeared in "High Performance Cluster Computing: Architectures and Systems", Prentice Hall, USA, 1999 (see <http://www.dgs.monash.edu.au/~rajkumar/cluster/>)

⁺ Networks of Workstations (NOW), Clusters of Workstations (COW) and Workstation Clusters are synonymous.

applications. Examples of these standards are the message passing library MPI¹ and data-parallel language HPF². In this context, standardisation enables applications to be developed, tested and even run on NOW and then at a later stage to be ported, with little modification, onto dedicated parallel platforms where CPU-time is accounted and charged for. The following list highlights some of the reasons for preference of NOW over specialised parallel computers^{3,4}:

- Individual workstations are becoming increasingly powerful.
- The communications bandwidth between workstations is increasing as new networking technologies and protocols are implemented in a LAN.
- Workstation clusters are easier to integrate into existing networks than special parallel computers.
- Typical low user utilisation of personal workstations.
- The development tools for workstations are mature than the contrasting proprietary solutions for parallel computers – mainly due to the non-standard nature of many parallel systems.
- Workstation clusters are a cheap and readily available alternative to specialised high performance computing platforms.
- Clusters can be easily grown – node capability can be easily extended (e.g. by adding memory blocks) and additional nodes can be added easily to extend cluster capability.

A Cluster can provide similar or better performance and reliability than traditional mainframes or supercomputers. Also, if designed correctly, better fault-tolerance at a much lower hardware cost can be achieved. Many applications have successfully used clusters including computationally intensive, distributed databases and telecommunications applications, to name but a few. In this paper, we focus on components, tools, techniques, and methodologies involved in using clusters for high-performance or parallel computing.

2. COMMODITY COMPONENTS FOR CLUSTERS

The continuous improvement of workstation and network performance and the availability of standardized message passing APIs are paving the way for usage of cluster-based parallel computing. The trends in hardware and software

¹ Snir, Otto, Huss-Lederman, D. Walker, and J. Dongarra, *MPI The Complete Reference*, MIT Press; 1996

² C. Koelbel, et al., *The High Performance Fortran Handbook*, The MIT Press, 1994.

³ M. Baker, et.al. *Review of Cluster Management Software*, NHSE Review, May 1996 – <http://www.nhse.org>

⁴ T. Anderson, D. Culler, and D. Patterson. *A Case for NOW (Network of Workstations)*. IEEE Micro, 15(1):54-64, Feb. 95.

technologies are discussed in later part of this paper. In this section, we discuss some of the hardware and software components commonly used while building clusters and its nodes.

2.1 Processors

Intel processors are those most commonly used in PC-based computers. The current generation Intel x86 processor family includes the Pentium Pro and II. These processors, whilst not in the high band of performance, match the performance of medium level workstation processors⁵. In the high performance band, the Pentium Pro shows a very strong integer performance, beating Sun's UltraSPARC at the same clock speed, however floating point performance is much lower. The Pentium II Xeon⁶, like the newer Pentium II's, uses a 100MHz memory bus. It is available with a choice of 512KB to 2MB of L2 cache, and the cache is clocked at the same speed as the CPU, overcoming the L2 cache size and performance issues of the plain Pentium II. The accompanying 450NX chipset for the Xeon supports 64-bit PCI busses, which could improve communication performance on Gigabit-performance interconnects.

Other popular processors include x86 variants (AMD x86, Cyrix x86), Alpha, PowerPC, SPARC. Computer systems based on these processors have also been used as cluster nodes, for instance, Berkeley NOW uses Sun's SPARC family processors based workstations as cluster nodes.

2.2 Memory and Cache

Originally PCs had 640 KBytes of memory usually "hardwired" onto the motherboard. Typically today a PC will be delivered with between 32 and 64 MBytes installed in slots with each slot holding a memory module. Computer systems can use various types of memory. These include Extended Data Out (EDO) which allows the next access to begin while the previous data is still being read and fast page which allows multiple adjacent accesses to be made more efficiently. Access to DRAM is extremely slow compared to the speed of the processor, taking up to orders of magnitude more time than a single CPU clock cycle. Caches are used to keep recently used blocks of memory for very fast access. However, the very fast memory used for cache is expensive and cache control circuitry becomes more complex as the size of the cache grows.

⁵ The Standard Performance Evaluation Corporation (SPEC) - <http://open.speclab.org>

⁶ Intel's Pentium II Xeon Processor - <http://www.tomshardware.com/xeon.html>

2.3 Disk IO

There are two disk interfaces commonly used within the PC. The first is IDE, although on early generations of the PC it was usually found on a daughter card, it is now often found built into Pentium mother boards in the Enhanced IDE (EIDE) form⁷. This is a 16-bit wide interface with a peak transfer rate of 5 MBytes/s transfer rates. The largest IDE drive commonly available is 10 Gbytes with a price of about \$35 per GByte. An IDE interface allows two devices to be connected, and there are typically only one or two interfaces within a PC. CD-ROM and tape drives are also available with IDE interfaces.

The other disk interface found within PCs is SCSI⁸. The Fast and Wide SCSI interface is capable of a peak rate of 10 MBytes/s in asynchronous and 20 MBytes/s in synchronous modes. A SCSI interface may contain up to 8 devices. For similar sizes, SCSI disks tend to outperform IDE drives, however IDE has a significant price advantage (as much as two to one). SCSI disks are however available with much larger capacities, up to about 18 GBytes is common and these retail at about \$55 per GByte.

2.4 Cluster Interconnects

The nodes in a cluster communicate over high-speed network using a standard networking protocol such as TCP/IP or a low-level protocol such as Active Messages. In most facilities it is likely that the interconnection will be via standard Ethernet. In terms of performance, (latency and bandwidth) this technology is showing its age. However, Ethernet is a cheap and easy way to provide file and printer sharing. A single Ethernet connection cannot be used seriously as the basis for cluster-based computing, its bandwidth and latency are not balanced compared to the computational power of the workstations now available. Typically one would expect the cluster inter-connect bandwidth to exceed of 10 MBytes/s and have message latencies of less than 100 μ s. A number of high performance network technologies are available in the marketplace, in this section we discuss a few of them.

2.4.1 Ethernet, Fast Ethernet, and Gigabit Ethernet

Standard Ethernet has become almost synonymous with workstation networking. This technology is in widespread usage, both in the academic and commercial sectors. However, its 10 Mbps bandwidth is no longer sufficient for use in environments where users are transferring large data quantities or there are high traffic densities. An improved version,

⁷ *The Enhanced IDE FAQ*, J. Wehman and P. Herweijer.

⁸ SCSI-1: Doc # X3.131-1986, ANSI, 1430 Broadway, NY, USA

commonly known as Fast Ethernet, provides 100 Mbps bandwidth and has been designed to provide an upgrade path for an existing Ethernet installations. Standard and Fast Ethernet cannot co-exist on a particular cable, but each uses the same cable type. When an installation is hub-based and uses twisted-pair it is possible to upgrade the hub to one, which supports both standards, and replace the Ethernet cards in only those machines where it is believed to be necessary.

Now, the-state-of-the-art Ethernet is Gigabit Ethernet⁺ and its attraction is largely due to two key characteristics. First, it preserves Ethernet's simplicity while enabling a smooth migration to gigabit-per-second (Gbps) speeds. Second, it delivers very high bandwidth to aggregate multiple Fast Ethernet segments and to support high-speed server connections, switched intra-building backbones, interswitch links and high-speed workgroup networks.

2.4.2 Asynchronous Transfer Mode (ATM)

ATM⁹ is a switched virtual-circuit technology and it was originally developed for the telecommunications industry. It is embodied within a set of protocols and standards defined by the International Telecommunications Union. The international ATM Forum, a non-profit organisation continues this work. Unlike some other networking technologies, ATM is intended to be used for both LAN and WAN, presenting a unified approach to both. ATM is based around small fixed sized data packets termed cells. It is designed to allow cells to be transferred using a number of different media such as both copper wire and fibre optic. This hardware variety also results in a number of different inter-connect performance levels.

When first introduced ATM used optical fibre as the link technology. However, in desktop environments this is undesirable. The two most common cabling technologies found in a desktop environment are telephone style cables (CAT-3) and a better quality cable (CAT-5). CAT-5 can be used with ATM allowing upgrades of existing networks without replacing cabling.

2.4.3 Scalable Coherent Interface (SCI)

SCI¹⁰ is an IEEE 1596 standard aimed at providing a low-latency shared memory across cluster. SCI is the modern equivalent of a Processor-Memory-I/O bus and LAN, combined. It is designed to support distributed multiprocessing with high bandwidth and low latency. It provides a scalable architecture that

⁺ "Gigabit Ethernet is Ethernet, only faster!"

⁹ *ATM User-Network Interface Specification*, ATM Forum/Prentice Hall, September 1993, ISBN 0-01-225863-3.

¹⁰ *Scalable Coherent Interface (SCI)*, IEEE 1596-1992, August 1993

allows large systems to be built out of many inexpensive mass-produced components¹¹. SCI is a point-to-point architecture with directory-based cache coherence. It can reduce the delay of inter-processor communications even when compared to the newest and best technologies currently available, such as Fibre Channel and ATM. SCI achieves this by eliminating the need for run-time layers of software protocol-paradigm translation. A remote communication in SCI takes place as just a part of a simple load or store process in a processor. Typically, a remote address results in a cache miss. This in turn causes the cache controller to address remote memory via SCI to get the data. The data is fetched to cache with a delay in the order of a few microseconds and then the processor continues execution.

2.4.4 Myrinet

Myrinet is a 1.28 Gbps full duplex LAN supplied by Myricom¹². It is a proprietary, high performance interconnect. Myrinet uses low latency cut-through routing switches, which is able to offer fault tolerance by automatic mapping of the network configuration, this also simplifies setting up the network. Myrinet supports both Linux and NT. In addition to TCP/IP support, the MPICH implementation of MPI¹³ is provided as well number of other customer developed packages, which provide sub 10µs latencies.

Myrinet is relatively expensive when compared to Fast Ethernet, but has real advantages over it: very low-latency (5µs, one-way point-to-point), very high throughput, and a programmable on-board processor allowing for greater flexibility. It can saturate the effective bandwidth of a PCI bus at almost 120 Mbytes/s with 4Kbytes packets¹⁴. The cost of Myrinet-LAN components, including the cables and switches, is in the range of \$1,500 per host. Also, switches with more than 16 ports are unavailable, so scaling can be messy, although switch chaining is used to construct larger Myrinet clusters.

2.5 Operating Systems



OS technology has matured and most systems can be easily extended to provide extra functionality without modifying the underlying structure. Unix and its variants (such as Sun Solaris, IBM's AIX, etc.) are popularly used to power workstations cluster nodes. In this section, we discuss two popular operating systems that are used to power nodes of Cluster of PCs (CoPs).

¹¹ SCI Association - <http://www.SCIZZL.com/>

¹² Myricom, Inc. - <http://www.myri.com/>

¹³ *MPI-FM: MPI for Fast Messages* - <http://www-csag.cs.uiuc.edu/projects/comm/mpi-fm.html>

¹⁴ N. Boden, et. al., *Myrinet - A Gigabit-per-Second Local-Area Network*, IEEE Micro, Vol. 15, No.1, February 1995.



A Survival Handbook for Multi-Platform Users, Developers, and Managers

MVS and UNIX covers the three major UNIX platforms, it compares and contrasts operating systems. More importantly, it's the best source of guidance in making the difficult choice between integration and migration. With over 600 pages, this handbook is a true survival guide for any mainframe professional whose environments and applications are moving to UNIX.

MVS and UNIX Richard J Bambara 0.07.006663.9 RRP\$143.95 AUUG\$115.00

Performance Measurement, Fine Tuning, and Capacity Planning for Releases 2.5.1 and 2.6

Solaris Performance Administration is the most comprehensive, useful, and sophisticated reference on Solaris 2.x ever published, this book gives you polished, professional tools for tweaking peak performance from both software and hardware and staying on top of capacity needs. Useful to both Solaris newcomers and seasoned pros, this guide features important background theory on queuing models and performance metrics, essential to assessment. It's packed with command output and code samples, plus technical and strategic information on everything from kernel architecture to network configuration and memory management – all with the benefit of tried-and-tested Solaris 2.x experience.

Solaris Performance Administration H Frank Cervone 0.07.011768.3 RRP\$104.95 AUUG\$83.95

UNIX

Save
20%

COMPUTING MCGRAW-HILL

Mail or fax your order to (02) 9417 7003 today for faster processing.

Please print clearly:

Name: _____

Phone:Bus/Home: () _____

Address: _____

Postcode: _____

Payment details: (Please tick)

I enclose payment now by: Cheque Bankcard American Express Mastercard Visa Diners Club

Card number:

Expiry Date: _____

Signature: _____

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

(Please note: If you pay by credit card, your credit card will not be debited until 21 days after your invoice date).

Qty	ISBN	TITLE	ARRP \$	AUUG Price \$
	0.07.006663.9	MVS and UNIX – Bambara s/c	143.95	115.00
	0.07.011768.3	Solaris Performance Administration – Cervone s/c	104.95	83.95
RRP=Recommended Retail Price. Prices and availability are subject to change without notice			Add \$6.00 Postage & Handling	
			TOTAL	
			6.00	

Please make sure you have enclosed payment with your order. If you fax your order please do not post your order also.



McGraw-Hill Book Company Australia Pty Ltd A.C.N. 000 473 674
4 Barcoo Street, Roseville NSW 2069. Tel: (02) 9417 4288 Fax: (02) 9417 7003
e-mail: cservice@mcgraw-hill.com.au

AUUG199

2.5.1 LINUX

Linux¹⁵ is a UNIX like OS, which was initially developed by Linus Torvalds, a Finnish undergraduate student in 1991-92. The original releases of Linux relied heavily on the Minix OS, however, the efforts of a number of collaborating programmers has resulted in the development and implementation of a robust and reliable, POSIX compliant, OS.

Although initially developed by a single author, there are now a large number of authors involved in the development of Linux. One major advantage of this distributed development has been that there is a wide range of software tools, libraries and utilities available. This is due to the fact that any capable programmer has access to the OS source and can implement the feature that they wish. Linux quality control is maintained by only allowing kernel releases from a single point. The availability via the Internet has led to fast feedback about bugs and other problems. The following are some of advantages of using Linux:

- It runs on inexpensive x86-based platforms.
- It offers the power and flexibility of UNIX.
- Linux is readily available on the Internet and can be download for free.
- It is easy to fix bugs and improve system performance.
- Users can develop or fine-tune hardware drivers and this can easily be made available to other users.

Linux provides the features typically found in UNIX implementations such as: pre-emptive multi-tasking, demand-paged virtual memory, multi-user and multi-processor support. Most applications written for UNIX will require little more than a re-compile. In addition to the Linux kernel, a large amount of application/systems software is also freely available, including GNU software and XFree86, a public domain X server.

2.5.2 Microsoft NT

Microsoft Corp.¹⁶ is the dominant provider of software in the personal computing market place. Microsoft provides two basic OSs: Windows 98 and Windows NT 4 (soon to become Windows NT 5)¹⁷. NT and Windows 98/95 had approximately 66% of desktop OS market share in 1996¹⁸. Windows NT is a 32-bit¹⁹, preemptive, multi-tasking, multi-user OS and supports SMP platform. NT is fault tolerant - each application operates in its own virtual memory address space. Unlike earlier version of Windows (such

¹⁵ [Linux Meta-FAQ](http://sunsite.unc.edu/mdw/linux.html) - <http://sunsite.unc.edu/mdw/linux.html>

¹⁶ [Microsoft Corporation](http://www.microsoft.com) - <http://www.microsoft.com>

¹⁷ A. Watts, *High Command*, PC Direct, December 97

¹⁸ [The Microsoft Market Share](http://newsport.sfsu.edu/ms/markets.htm) - <http://newsport.sfsu.edu/ms/markets.htm>

¹⁹ H. Custer, *Inside Windows NT*, Microsoft Press, 1993, ISBN 1-55615-481-X

as Windows for Workgroups and Windows 95), NT is a complete OS, and not an addition to DOS. NT supports most popular CPUs architectures. NT has an object-based security model and its own special file system (NTFS) that allows permissions to be set on a file and directory basis. Packaged with NT are several built-in networking protocols, such as - IPX/SPX, TCP/IP, and NetBEUI) and APIs, such as - NetBIOS, DCE RPC, and Windows Sockets.

3. PROGRAMMING TOOLS FOR HPC ON CLUSTERS

The availability of standard programming tools and utilities on NOW has made clusters a practical proposition for their usage in parallel computing. In this section we discuss a few popular tools that are available on clusters.

3.1 Message Passing Systems (PVM/MPI)

Message passing libraries allow efficient parallel programs to be written for distributed memory systems. These libraries provide routines to initiate and configure the messaging environment as well as sending and receiving packets of data. Currently, the two most popular high-level message-passing systems for scientific and engineering application are the Parallel Virtual Machine²⁰ (PVM) and the Message Passing Interface²¹ (MPI).

PVM is both an environment and a message-passing library, which can be used to run applications on systems ranging from high-end supercomputers through to NOW. The MPI standard is the amalgamation of what were considered the best aspects of the most popular message-passing systems at the time of its conception. The standard only defines a message passing library and leaves, amongst other things, the initialisation and control of processes to individual developers to define. MPI is fast becoming a *de facto* standard. MPI and PVM libraries are available for Fortran 77 and 90, ANSI C and C++. There also exist interfaces to other languages - one such example is Java²². Commercial and public-domain²³ implementations of MPI/PVM exist and they run on a range of platforms from tightly coupled systems to NOWs.

²⁰ A. Beguelin, et. al., *The PVM project, Technical Report*, Oak Ridge National Laboratory, February 1993.

²¹ MPI Forum, *MPI: A Message-Passing Interface Standard*, University of Tennessee, Report CS-94-230, May 1994

²² [mpiJava Wrapper](http://www.npac.syr.edu/projects/prpc/m_piJava/) - http://www.npac.syr.edu/projects/prpc/m_piJava/, August 1998.

²³ M.A. Baker and G.C. Fox, *MPI on NT: A Preliminary Evaluation of the Available Environments*, 12th IPPS & 9th SPDP, LNCS, Jose Rolim (Ed.), Parallel and Distributed Computing, Springer Verlag, Germany. ISBN 3-540 64359-1.

3.2 Distributed Shared Memory (DSM) Systems

The most efficient, and widely used, programming paradigm on distributed memory systems is message passing. A problem with this paradigm is that it is complex and difficult to program compared to shared-memory programming systems. Shared memory systems offer a simple and general programming model, but they suffer from scalability. An alternate cost effective solution is to build a DSM system on distributed memory system, which exhibits simple and general programming model and scalability of distributed memory systems.

DSM systems can be implemented by using software or hardware techniques. The characteristics of software DSM systems are: they are usually built as a separate layer on top of the communications interface; they take full advantage of the application characteristics; virtual pages, objects, and language types are units of sharing. Software DSM can be implemented either solely by runtime, compile time, or combined approaches. Two representative software DSM systems are TreadMarks²⁴ and Linda²⁵. The characteristics of hardware DSM systems are: better performance (much faster than software DSM), no burden on user and software layers, fine granularity of sharing; extensions of the cache coherence schemes, and increased hardware complexity. Two examples of hardware DSM systems are DASH²⁶ and Merlin²⁷.

3.3 Parallel Debuggers and Profilers

To develop correct and efficient high performance applications it is highly desirable to have some form of easy-to-use parallel debugger and performance profiling tools. Most vendors of HPC systems provide some form of debugger and performance analyzer for their platforms. Ideally, these tools should be able to work in a heterogeneous environment. Thus making it possible to develop and implement a parallel application on, say a NOW, and then actually do production runs on a dedicated HPC platform, such as the Cray T3E.

3.3.1 Debuggers

The number of parallel debuggers that are

²⁴ *TreadMarks*
<http://www.cs.rice.edu/~willy/TreadMarks/overview.html>

²⁵ N. Carriero and D. Gelernter, *Linda in context*, Communications of the ACM, 32(4): 444-458, April 1989.

²⁶ Lenoski, D. Laudon J, et al., *The Stanford DASH Multiprocessor*, IEEE Computer, Vol. 25, No. 3, March 1992.

²⁷ C. Mapples, Li Wittie, *Merlin: A Superglue for Multiprocessor Systems*, CAMPCON'90, March 1990.

capable of being used in a cross-platform, heterogeneous, development environment is very limited²⁸. Hence, in 1996 an effort was begun to define a cross-platform parallel debugging standard that defined the features and interface users wanted. The High Performance Debugging Forum (HPDF) was formed as a Parallel Tools Consortium²⁹ project. The forum has developed a HPD Version specification, which defines the functionality, semantics, and syntax for a command-line parallel debugger. Ideally, a parallel debugger should be capable of:

- Managing multiple processes, and multiple threads within a process.
- Displaying each process in its own window.
- Displaying source code, stack trace and stack frame for one or more processes.
- Diving into objects, subroutines and functions.
- Setting both source-level and machine-level breakpoints.
- Sharing breakpoints between groups of processes.
- Defining watch and evaluation points.
- Displaying arrays and array slices.
- Manipulation of code variables and constants.

TotalView is a commercial product from Dolphin Interconnect Solutions³⁰. It is currently the only widely available parallel debugger that supports multiple HPC platforms. TotalView supports most commonly used scientific languages (C, C++, F77/F90 and HPF), message passing libraries (PVM/MPI) and most common UNIX OS. Even though TotalView can run on multiple platforms, it can only be used in an homogeneous environments.

3.4 Performance Analysis Tools

The basic purpose of performance analysis tools is to help a programmer to understand the performance characteristics of a particular application. In particular analyze and locate parts of an application that exhibit poor performance and create program bottlenecks. Such tools are useful for understanding the behavior of normal sequential applications and can be enormously helpful when trying to analyze the performance characteristics of parallel applications.

Most performance monitoring tools consist of the some or all of the following components:

- A means of inserting instrumentation calls to the performance monitoring routines into the user's application.
- A runtime performance library that consists of a set of monitoring routines that measure and record various aspects of programs

²⁸ S. Browne, *Cross-Platform Parallel Debugging and performance Analysis Tools*, Procs. of the EuroPVM/MPI98 Workshop, Liverpool, September 1998.

²⁹ *Parallel Tools Consortium project*
<http://www.ptools.org/>

³⁰ *Dolphin Interconnect Solutions*
<http://www.dolphinics.no/>

performance.

- A set of tools that process and displays the performance data.

A particular issue with performance monitoring tools is the intrusiveness of the tracing calls and their impact on the applications performance. It is very important to note that instrumentation affect the performance characteristics of the parallel application and thus provide a false view of its performance behavior. The most commonly used tools for performance analysis of message passing programs are:

- AIMS³¹ - instrumentors, monitoring library, and analysis tools.
- MPE³² - logging library and Nupshot performance visualization tool.
- Pablo³³ - monitoring library and analysis tools.
- Paradyn³⁴ - dynamic instrumentation and run-time analysis tool.
- SvPablo³³ - integrated instrumentor, monitoring library, and analysis tool.
- VAMPIRtrace³⁵ - monitoring library and VAMPIR performance visualization tool.
- Dimemas³⁶ - a performance prediction tool for message passing programs.
- PARAVR³⁷ - a parallel program visualization and analysis tool.

3.5 Cluster Monitoring – System Administration Tools

Monitoring clusters is a challenging task that can be eased by tools which allow the entire cluster system to be observed at different levels using a GUI. There are many projects investigating system administration on clusters, including Berkeley NOW⁴¹, K-CAP³⁸, and PARMON³⁹. The Berkeley system administration tool gathers and stores data in a relational database. It uses a Java applet as primary interface to allow users to monitor the system

³¹ AIMS
<http://science.nas.nasa.gov/Software/AIMS>

³² MPE - <http://www.mcs.anl.gov/mpi/mpich>

³³ Pablo - <http://www-pablo.cs.uiuc.edu/Projects/Pablo/>

³⁴ Paradyn - <http://www.cs.wisc.edu/paradyn/>

³⁵ VAMPIR - <http://www.pallas.de/pages/vampir.htm>

³⁶ Dememas - <http://www.pallas.com/pages/dimemas.htm>

³⁷ PARAVR - <http://www.cepba.upc.es/paraver/>

³⁸ P. Uthayopas, et. al, *Interactive Management of Workstation Clusters Using World Wide Web*, Cluster Computing Conference-CCC '97, Proceedings, <http://www.mathcs.emory.edu/~ccc97/>, 1997.

³⁹ Rajkumar, Krishnamohan and Bindu, *PARMON: A Comprehensive Cluster Monitoring System*, AUUG'98 - Open Systems: The Common Thread, Sydney, Australia, 1998

via the Internet from their Web browser⁴⁰. K-CAP, uses the Web and Java Applets for connecting its management node to a predefined URL addresses in the cluster. PARMON allows monitoring of resources at component, node, and the entire system levels.

4. REPRESENTATIVE CLUSTER SYSTEMS

There are many projects investigating the development of supercomputing class system from COTS components. We briefly describe the following popular efforts:

- Network of Workstations (NOW) project at University of California, Berkeley.
- High Performance Virtual Machine (HPVM) project at University of Illinois at Urbana-Champaign.
- Beowulf Project at the Goddard Space Flight Center, NASA;

4.1 Network Of Workstations (NOW) Project

The Berkeley NOW project⁴¹ demonstrates how a successful large-scale parallel computing systems can be put together with volume produced commercial workstations and the latest commodity switch-based network components. To attain the goal of combining distributed workstations into a single system the NOW project included research and development into network interface hardware, fast communication protocols, distributed file systems, distributed scheduling and job control. The Berkeley system consists of the following:

Inter-processor Communications - Active Messages (AM) is the basic communications primitives in Berkeley NOW. It generalises previous AM interfaces to support a broader spectrum of applications such as client/server programs, file systems, operating systems, as well continuing support for parallel programs. The AM communication is essentially a simplified remote procedure call that can be implemented efficiently on a wide range of hardware. NOW includes a collection of low-latency, parallel communication primitives: Berkeley Sockets, Fast Sockets, shared address space parallel C (Split-C), MPI and a version of HPF.

Global Layer Unix (GLUnix) - is an OS layer designed to provide transparent remote execution, support for interactive parallel and sequential jobs, load balancing, and backward compatibility for existing application binaries. GLUnix is a multi-user system implemented at the user-level so that it can be easily ported to a

⁴⁰ E. Anderson and D. Patterson, *Extensible, Scalable Monitoring for Clusters of Computers*. Proceedings of the 11th Systems Administration Conference (LISA '97), October 26-31, 1997, San Diego, California, USA.

⁴¹ Berkeley NOW - <http://now.cs.berkeley.edu/>

number of different platforms. GLUnix aims to provide: A cluster-wide namespace and uses Network PIDs (NPIDs) and Virtual Node Numbers (VNNs). NPIDs are globally unique process identifiers for both sequential and parallel programs throughout the system. VNNs are used to facilitate communications among processes of a parallel program. A suite of user tools for interacting and manipulating NPIDs and VNNs, equivalent to UNIX run, kill, make and tcsh. A GLUnix API allows interaction with NPIDs and VNNs.

Network RAM - is a means to utilise free resources on idle machines as a paging device for busy machines. The designed system is serverless, and any machine can be a server when it is idle, or a client when it needs more memory than physically available. Two prototype systems have been developed. One of these uses custom Solaris segment drivers to implement an external user-level pager, which exchanges pages with remote page daemons. The other provides similar operations on similarly mapped regions using signals.

xFs File System - is a serverless, distributed file system, which attempts to have low latency, high bandwidth access to file system data by distributing the functionality of the server among the clients. The typical duties of a server include maintaining cache coherence, locating data, and servicing disk requests. The function of locating data in xFS is distributed by having each client responsible for servicing requests on a subset of the files. File data is striped across multiple clients to provide high bandwidth.

4.2 The High Performance Virtual Machine (HPVM) Project

The goal of the HPVM project⁴² is to deliver supercomputer performance on low-cost COTS systems. HPVM also aims to hide the complexities of distributed system behind clean interface. The HPVM project provides software that enables high-performance computing on clusters of PCs and workstations. The HPVM architecture consists of a number of software components with high-level APIs, such as MPI, SHMEM and Global Arrays, that allows HPVM clusters to be competitive with dedicated MPP systems.

The HPVM project aims to address the following challenges:

- Delivering high-performance communication to standard, high-level APIs.
- Co-ordinating scheduling and resource management.
- Managing heterogeneity.

A critical part of HPVM is a high-bandwidth and low-latency communications protocol known as Fast Messages (FM), which is based on Berkeley AM. Unlike other messaging layers, FM is not the

surface API, but the underlying semantics. FM contains functions for sending long and short messages and for extracting messages from the network. The services provided by FM guarantees and controls the memory hierarchy that FM provides to software built with FM. FM also guarantees reliable and ordered packet delivery as well as control over the scheduling of communication work.

The FM interface was originally developed on a Cray T3D and a cluster of SPARCstations connected by Myrinet hardware. FM has a low-level software interface that delivers hardware communication performance; however, higher-level layers interface offer greater functionality, application portability and ease of use.

4.3 The Beowulf Project

The Beowulf project⁴³ aim was to investigate the potential of PC clusters for performing computational tasks. Beowulf refers to a Pile-of-PCs (PoPC) to describe a loose ensemble or cluster of PCs, which is similar to COW/NOW. PoPC emphasizes on the use of mass-market commodity components, dedicated processors (rather than stealing cycles from idle workstations), and the usage of a private communications network. An overall goal of Beowulf is to achieve the "best" overall system cost/performance ratio for the cluster.

System Software

The collection of software tools being developed and evolving within the Beowulf project is known as "Grendel". These tools are for resource management and to support distributed applications. The Beowulf distribution includes several programming environments and development libraries as separate packages. These include PVM, MPI, and BSP, as well as, SYSV- style IPC and pthreads.

The communication between processors in Beowulf is through TCP/IP over the Ethernet internal to cluster. The performance of inter-processor communications is, therefore, limited by the performance characteristics of the Ethernet and the system software managing for message passing. Beowulf has been used to explore the feasibility of employing multiple Ethernet networks in parallel to satisfy the internal data transfer bandwidths required. Each Beowulf workstation has user-transparent access to multiple parallel Ethernet networks. This architecture was achieved by "channel bonding" techniques implemented as a number of enhancements to the Linux kernel. The Beowulf project has shown that up to three networks can be ganged together to obtain significant throughput, thus validating their usage of the channel bonding technique. The

⁴² *HPVM* - <http://www-csag.cs.uiuc.edu/projects/clusters.html>

⁴³ *The Beowulf Project* - <http://www.beowulf.org>

usage of new network technologies, such as Fast Ethernet, will ensure even better inter-processor communications performance.

In the interests of presenting a uniform system image to both users and applications, Beowulf has extended the Linux kernel to allow a loose ensemble of nodes to participate in a number of global namespaces. In a distributed scheme it is often convenient for processes to have a PID that is unique across an entire cluster, spanning several kernels. Beowulf implements two Global Process ID (GPID) schemes. The first is independent of external libraries. The second, GPID-PVM, is designed to be compatible with PVM Task ID format and use PVM as its signal transport. While the GPID extension is sufficient for cluster-wide control and signaling of processes, it is of little use with a global view of the processes. To this end, the Beowulf project is developing a mechanism that allows unmodified versions of standard UNIX utilities (e.g., ps) to work across a cluster.

4.4 A Comparison of the Three Cluster Environments

The cluster projects described in this paper share a common goal of attempting to provide a unified resource out of inter-connected PCs or workstations. Each system claims that it is capable of providing supercomputing resources from COTS components. Each project provides these resources in different ways, both in terms how the hardware is connected together and the way the system software and tools provide the services for parallel applications.

Table 1 shows the key hardware and software components that each system uses. Beowulf and HPVM are capable of using any PC, whereas Berkeley NOW functions on platforms where Solaris is available. Berkeley NOW and HPVM use Myrinet with a fast, low-level communications protocol (AM and FM).

System	Hardware		Software	
	Platform	Communications	OS	Other
Beowulf	PCs	Multiple Ethernet with TCP/IP	Linux + Grendel	MPI/PVM, Sockets and HPF
Berkeley NOW	Solaris-based PCs and workstations	Myrinet with Active Messages.	Solaris + GLUuix + xFS	AM, PVM/MPI, HPF and Split-C
HPVM	PCs	Myrinet with Fast Messages.	NT or Linux + connection and global resource manager + LSF	Java front-end, FM, Sockets, Global Arrays, SHMEM and MPI.

Table 1 - Cluster Systems Comparison

Each system consists of some middleware interfaced into the OS kernel, which is used to provide a globalization layer, or unified view of the distributed cluster resources. Berkeley NOW use the Solaris OS, whereas Beowulf uses Linux with a modified kernel and HPVM is available for both Linux and Windows NT. All three systems provide a rich variety of tools and utilities commonly used to develop, test and run parallel applications. These include various high-level APIs for message passing and shared-memory programming.

5. SUMMARY AND CONCLUSIONS

In this paper we have briefly discussed the different hardware and software components that are commonly used in the current generation cluster-based systems. We have also described three state-of-the-art projects that are using subtly different approaches ranging from an all COTS approach through to a mixture of technologies. In this section we summarise our findings, and then make a few comments about possible future trends.

5.1 Hardware and Software Trends

In the last five years several important advances have taken place and prominent among these are:

- A network performance increase of 10 fold using 100BaseT Ethernet with full duplex support
- The availability of switched network circuits, including full crossbar switches for proprietary network technologies such as Myrinet
- Workstation performance has improved significantly.
- Improvement of microprocessor performance has led to availability of desktop PCs with performance of low-end workstations, but at significantly lower cost.
- The availability of fast, functional and stable OSs (Linux) for PCs, with source code access.
- The performance gap between supercomputer and commodity-based clusters is closing rapidly.

- Parallel supercomputers are now equipped with COTS components, especially microprocessors (Cray T3E - DEC Alpha), whereas earlier systems had custom components.

Increasing usage of SMP nodes with two to four processors.

Culler and Singh⁴⁴ quantify a number of hardware trends. Foremost of these is the design and manufacture of microprocessors. A basic advance is the decrease in feature size which enables circuits to become either faster or lower in power consumption. In conjunction with this is the growing die size that can be manufactured. These factors mean that:

- The average number of transistors on a chip is growing by about 40% per annum.
- The clock frequency growth rate is about 30% per annum.

It is anticipated that by early 2000 there will be 700 MHz processors with about 100 million transistors. There is a similar story for storage but the divergence between memory capacity and speed is more pronounced. Memory capacity increased by three orders of magnitude between 1980 and 1995, yet its speed has only doubled.

It is anticipated that Gigabit DRAM will be available in early 2000, but the gap to processor speed is getting greater all the time. The problem is that memories are getting larger whilst processors are getting faster. So getting access to data in memory is becoming a bottleneck. One method of overcoming this bottleneck is to configure the DRAM in banks and then transfer data from these banks in parallel. In addition, multi-level memory hierarchies organised as caches make memory access more effective, but their design is complicated. The access bottleneck also applies to disk access, which can also take advantage to parallel disks and caches.

The ratio between the cost and performance of network interconnects is falling rapidly. New network technologies are being used successfully by many projects, but no single network technology has yet emerged as a clear winner. Myrinet is not a commodity product and costs a lot more than Ethernet, but it has the advantages of low-latency, high bandwidth and a programmable interface. SCI has been used to build distributed shared memory system, but lacks scalability. ATM is used in clusters that are mainly used for multimedia processing.

Linux has become a popular alternative to commercial OSs due to being freely available and having superior performance to Microsoft offerings. Linux currently has more than 7 millions users worldwide and has become the researchers OS of choice. NT also has a large installed user base. NT 5, which is due to be released in 1999, has a number of new features, including a thinner and faster TCP/IP stack. NT

systems for parallel computing are in a similar situation to UNIX workstation 5/7 years ago. It seems that it will only be a matter of time before NT catches up with UNIX-based platforms in terms of functionality and performance.

5.2 Cluster Technology Trends

We have discussed a number of cluster projects within this paper. These range from those which are propriety-based (Solaris) through to a totally commodity system (Beowulf). HPVMM can be considered as a hybrid-system using commodity computers and specialised network interfaces. It should be noted that the projects detailed in this paper are a few of the most popular and well known, rather than an exhaustive list of all those available. All the projects discussed claim to consist of commodity components. This is true, however, one could argue that true commodity technologies would be those that are pervasive at most academic or industrial sites. If this were the case, then true commodity would mean PCs running Windows 95 with standard 10 Mbps Ethernet. However, when considering parallel applications with demanding computational and network needs, this type of low-end cluster would be incapable of providing the resources needed.

Each of the projects discussed tries to overcome the bottlenecks that using cluster-based systems for demanding parallel applications in a slightly different ways. Without fail, however, the main bottleneck is not the computational resource (be it a PC or UNIX workstation), rather it is the provision of a low-latency, high-bandwidth inter-connect and an efficient low-level communications protocol to provide high-level APIs.

The Beowulf project explores the usage of multiple standard Ethernet cards to overcome the communications bottleneck. Whereas Berkeley NOW and HPVN use programmable Myrinet cards and AM/FM communications protocols. The choice of what is the best solution cannot just be based on performance, the cost per node to provide the NIC should also be considered. For example, a standard Ethernet card costs less than \$100, whereas Myrinet cards cost in excess of \$1000 each. Another factor that must also be considered in this equation is the availability of Fast Ethernet and the advent of GigaBit Ethernet. It seems that Ethernet technologies are more likely to be main stream, mass produced and consequently cheaper than specialised network interfaces. As an aside, all the projects that have been discussed are in the vanguard of the cluster computing revolution and their research is helping the following army determine which are the best techniques and technologies to adopt.

⁴⁴ D. E. Culler, J. P. Singh, A. Gupta, *Parallel Computer Architecture: A Hardware/Software Approach*, Morgan Kaufmann Publishers, USA, 1998.


5.3 Future Cluster Technologies

Emerging hardware technologies along with maturing software resources mean that cluster-based systems are rapidly closing the performance gap with dedicated parallel computing platforms. Cluster systems that scavenge idle cycles from PCs and workstations will continue to use whatever hardware and software components available on public workstations. Clusters dedicated to high performance applications will continue to evolve as new and more powerful computers and network interfaces become available in the market place.

It is likely that individual cluster nodes will comprise of multiple processors. Currently two and four processor PCs and UNIX workstations are becoming common. Software that allows SMP nodes to be efficiently and effectively used by parallel applications will be developed and added to the OS kernel in the near future. It is likely that there will be widespread usage of Gigabit Ethernet and as such it will become the *de facto* standard for clusters. To reduce message passing latencies cluster software systems will by-pass the OS kernel, thus avoiding the need for expensive system calls, and exploit the usage of intelligent network cards. This can obviously be achieved using intelligent NICs, or alternatively using on-chip network interfaces such as those used by the new DEC Alpha 21364.

The choice of OS used on future clusters will be determined by its ability to provide a rich set of development tools and utilities as well as the provision of robust and reliable services. UNIX-based OSs are likely to be most popular, but the steady improvement and acceptance of Window NT will mean that it will be not far behind.

❖



AUREMA

Exciting Product Development Opportunities in Sydney

- Global business: AUS, US, UK
- Australian owned and operated
- Innovative work environment
- State of the art technology

Aurema is delivering world class active resource management products on leading operating systems such as Solaris and IRIX. Our dedicated focus on resource management technology provides a total solution in the allocation and control of major system resources on UNIX servers.

We are gearing up to deploy a suite of new and exciting products in the year 2000 on SCO UNIX, Linux and Windows 2000. Our goal is paramount ease of use, reliability and functionality - and we need the right people to help us deliver these products to the world market.

Project Manager

This senior role is for an experienced technical project manager to lead a group of software development teams developing both system and application level software predominantly on UNIX.

Test Team Leader

Contribute to the design and implementation of the test strategy, interacting closely with project leads, management of the test schedule/planning and test team, and verification of initial test reports.

Test Engineer

This challenging role involves testing of operating system enhancements and graphical user interfaces on a variety of UNIX platforms. You must have developed tests using C or C++ and UNIX scripting languages.

Senior Technical Writer

A senior technical writer is sought who has a very strong technical background and experience in the entire documentation life cycle. FrameMaker experience preferred.

System Software Engineer

A software engineer with at least four years experience is needed to design, develop and test components of our new product range. UNIX/C development experience essential.

Quality Engineer

You will assist the Quality Engineering Manager in the identification, development, implementation, training, and maintenance of the company Quality Management System.

WinNT Software Engineer

This role involves analysis, design, and prototyping of a Windows 2000 service. Good grounding in the Win32 API, Windows NT, and experience in developing Win32 applications on NT is essential, as are excellent communication skills.

Linux Software Engineer

This role involves analysis, design, and prototyping enhancements to Linux. Familiarity with Linux internals, an understanding of the open source community, experience in developing a Linux kernel module is essential.

Aurema provides a **stimulating and flexible work environment** close to Central Station. You will be working alongside some of the world's best systems software and applications developers. Aurema provides **regular training and development opportunities**. There are also opportunities for short term overseas assignments. A **very competitive salary** will be awarded based upon your experience.

Applications must be made in writing to the Vice President, Engineering, Aurema Pty Ltd. E-mail: info@arema.com Post: PO Box 305, Strawberry Hills, NSW 2012 Fax: (02) 9699 9174 For further information call (02) 9698 2322.

Visit our web site for further details: <http://www.aurema.com>

NEW!! Java technology seminars from Sun

Sun™ Educational Services has developed a series of seminars designed to allow programmers and developers to quickly master the latest Java™ technologies and maintain a competitive advantage in the industry:

Enterprise JavaBeans Programming

(seminar code: SEM-SL-351)

2 day seminar
\$1,200/person

(special price for AUUG members = \$1,080/person)

Seminar dates

Sydney: 18 October, 6 December
Melbourne: 4 October, 13 December

The Enterprise JavaBeans™ Programming seminar provides students with the framework for developing an enterprise application using Enterprise JavaBeans (EJB) technology. At the core of the Enterprise Java platform, EJB components simplify the development of flexible, multitiered, transactional applications. Enterprise JavaBeans students will gain knowledge to create interoperable components that are used with scalable and extensible application servers.

Who can benefit

Java architects, systems analysts, enterprise application tool developers, software integrators for using application servers and application developers for Business Enterprise transactional software.

Prerequisites

To gain the most from this seminar, you should be:

- Experienced with the Java programming language
- Experienced with object-oriented design and analysis
- Familiar with distributed programming (multitier architecture)
- Familiar with relational or object database programming
- Familiar with transactions
- Able to create multitier Java Application Solutions
- Familiar with component technology

Skills gained

Upon completion of this seminar, you should be able to:

- Describe the framework for developing an Enterprise application which adheres to the EJB 1.1 specifications
- List and describe the Enterprise APIs for the Java 2 Platform Standard Edition
- Describe CORBA basics
- Describe EJB Technology
- Access the EJB
- Describe miscellaneous EJB Topics and the future of the EJB specification
- Design the Session Bean and Entity Bean

Programming Distributed Services with Jini Technology

(seminar code: SEM-SL-370)

2 day seminar

\$1,200

(special price for AUUG members = \$1,080/person)

Seminar dates

Sydney: 21 October, 9 December

Melbourne: 7 October

This seminar provides students with the information needed to use Jini™ technology in real world applications. It covers the fundamentals of RMI, which is often used with Jini technology. The seminar describes the vision of Jini technology and discusses the concepts of spontaneous networking. It then discusses the key parts of the Jini technology programming model. These are leases, distributed events, transactions and the Jini Lookup service.

Who can benefit

Software engineers and architects who are competent in the Java programming language and are interested in programming distributed systems.

Prerequisites

To gain the most from this seminar, you should be:

- Able to design and implement software systems using the Java programming language
- Familiar with RMI and other distributed computer technology

Skills gained

Upon completion of this seminar, you should know:

- The vision behind Jini technology
- Distributed programming concepts
- RMI
- Leasing
- Distributed events
- JavaSpaces technology
- Transactions
- Jini Lookup and Discovery

Don't forget about the AUUG member DISCOUNTS!

10% off all instructor-led courses

15% off all CD-ROM courses

... as well as a **20% discount off all Sun Certification exams:**

To register for either of the Java technology seminars
or for information on any Sun Educational Services instructor-led courses,
CD-ROM courses or Sun Certification programs, phone:


1800 550 786 / 1800 550 SUN

or visit the Sun Educational Services web page:

<http://www.sun.com.au/services/education>

© 1999 Sun Microsystems, Inc., All rights reserved. Sun Microsystems, the Sun logo, Java, JavaBeans and Jini trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Priority code AV19101



Member Services Update: **New Online News Service**

EXCLUSIVE ON-LINE NEWS SERVICE PROPOSED FOR AUUG MEMBERS.

AUUG in partnership with IDG Communications, the world's foremost IT publisher, will be offering members a comprehensive UNIX information on-line news service. The free service will be available daily or weekly depending on your preference and adds great value to your existing AUUG membership.

Run by the ComputerWorld division of IDG, the service will pull content from IDG's global resources of more than 280 publications and from all its local magazines and services.

The content is specifically tailored to your information needs. AUUG will be defining the areas of interest that the tailored email service will be based upon. As a subscriber, you will receive a daily or weekly email with headlines and summaries of the stories pointing back to the full text of the article on a dedicated web site. A fully searchable archive of news and feature stories will provide a comprehensive and continuously growing information resource for AUUG members.

The service has strict privacy policies and your email address will only be used for this specific news report. You can subscribe or unsubscribe at any time.

This service brings AUUG closer to our members and we are pleased to be able to offer this to you.

To join the service send an email to auugnews@email.idg.com.au and include the phrase `join auugnews` in the body of the email.



Call for Papers/ Expression of Interest

Security is always a key interest of AUUG members. Recognising this, AUUG is holding the inaugural AUUG Security Symposium in Sydney during November 1999.

Where/When:

Tuesday 23 November 1999
Wesley Conference Centre
220 Pitt Street Sydney

Security is not something that can be looked at once, implemented and then forgotten - it must be part of the day to day operations. The nature of the challenge is that security threats present a moving target. Security must be a prime consideration in the design, implementation and operation of all environments. Security cannot be ignored and must be built in to new software and new installations.

The AUUG Security Symposium will focus on practical solutions to current problems. It will be a forum for professionals to explore what is happening in the sector, benefit from others' experiences, and acquire skills to deal with current and potential security problems.

Topics will include:

- Network security
- Systems security
- Security holes that are actively being exploited
- General principals for securing sites
- Software development in a secure framework
- What to do when you are told "just buy a firewall"

Who should attend?

- Network and Systems administrators
- Software developers

Timetable:

- Abstracts (around 100) words are due 1 October 1999.

Presenters will receive free registration.



If you have any questions, wish to register an Expression of Interest, or would like more information, please contact the AUUG Business on 1800 625 655 or email: busmgr@auug.org.au



My Home Network

Frank Crawford
frank@crawford.emu.id.au

Well, a new edition of AUUGN rolls around, and as promised here is the first official column of "My Home Network" (as opposed to the article in the previous edition).

What I intend to do with this column is to give more details about various items that compose my home network. I'll also report on any interesting feedback I receive about the column. So far I've received one reply, from Paul Pyyvaara. He has a network at his home (see I'm not the only nerd), although he has wired his house for UTP, even the garage!

I hope others will also respond with comments, corrections and stories. Hopefully you noticed a new email address for me:

frank@crawford.emu.id.au

It is email to my home network, which I'll discuss in a future column.

Anyway, onto the item I want to cover in this column, the auto-dialing daemon, 'diald'. It was interesting reading in the blurb for Microsoft Windows 98 Second Edition that it now includes an "Internet Connection" facility, which allows you to share your dialup connection with other systems in your network. Sounds like magic (at least to a Windows users :-)). Someone who has been using Unix for some time will recognise this as just setting up routing correctly, some form of NAT (Network Address Translation) and a daemon to do the dialing. In fact the Internet Connection side is really only the first two items (routing & NAT).

If Unix doesn't do routing right, then the whole of the Internet is in big trouble. The only interest here is that you often have to enable network forwarding, as it is generally disabled in the interest of security. For RedHat Linux, this just involves editing or adding a line in '/etc/sysconfig/network' to say:

```
"FORWARD_IPV4=true"
```

or in recent kernels (i.e. Linux 2.2 or 2.3) just add a line like:

```
`echo "1" > /proc/sys/net/ipv4/ip_forward'
```

to the startup scripts.

The issue of NAT is a bit more complicated, and is only needed if you wish to hide all your machines behind a single address (which would be the case for most home networks). Linux provides a facility called IP Masquerading. Much of this is compile time options, so I'm not going into it here, but if you need further details feel free to contact me. I should note that IP Masquerading isn't a full NAT

implementation (that is available elsewhere) but something that is specifically tailored to hiding a network behind a single IP address for dialup access. Along with this you will also need to look into the program 'ipchains' (or ipfwadm for Linux 2.0), to configure the masquerading.

So, this just leaves the automatic handling of demand dialing. This is exactly what diald does. To quote from the FAQ:

Diald is a daemon that does demand dialing for PPP and SLIP. The purpose of diald is to make it transparently appear that you have a permanent connection to a remote site. Diald sets up a "proxy" device which stands in for the physical connection to a remote site. It then monitors the proxy, waiting for packets to arrive. When interesting packets arrive it will attempt to establish the physical link to the remote site using either SLIP or PPP, and if it succeeds it will forward traffic from the proxy to the physical link. As well, diald will monitor traffic once the physical link is up, and when it has determined that the link is idle, the remote connection is terminated. The criteria for bringing the link up and taking it down are configurable at run time, and are based upon the type of traffic passing over the link.

The current version (or at least what I am using) 0.16, which is available as a RPM for RedHat Linux in the contrib section. For more information about 'diald' you can check the Diald home page:

<http://www.loonie.net/~eschenk/diald.html>

although it appears it will shortly be relocated, because there is a new maintainer.

Obviously, diald needs other support, and in particular, it needs SLIP in the kernel (as this is how it sets up the proxy device) and possibly PPP, if this is used for the link. In my case I also need 'pppd' for the connection and all associated software, e.g. 'chat' or 'expect'.

To set up the connection, you need to configure the phone number, end-point IP addresses (which can be dynamic) and routing details. In particular, if the link is to be the default route for your subnet. It can handle multiple destinations (although I haven't played with this feature yet) and setup 'proxyarp' connection so it looks like you are there. proxyarp is where an ARP entry is established for the remote systems IP address, so to other systems it looks like the remote system is on your local subnet. This only works if the IP address allocated for your PPP connection is valid on your network.

Along with 'diald', I've picked up RedHat contributed packages called 'diald-config' and 'diald-config-unmetered' (in noarch) which fits the configuration much more neatly into RedHat standard.

Finally, I've also picked up a useful control

program called `diald-top'. diald comes with `dctrl' which is an X11 based utility, but diald-top is a curses based utility which means I can use it from almost anywhere. It does have some bugs, and hasn't changed in nearly 2 years, but it seems to work fine. There isn't an RPM for it, but you can obtain the source from:

`sunsite.unc.edu:/pub/Linux/system/network/serial`

Just to finish off about diald, this package has been indispensable over the years, but has proved slightly more expensive than controlling dialups manually, because computers do tend to perform connections unexpectedly and at odd times. There are lots of options I haven't played with to handle these occasions, as I really haven't found it necessary.

As a postscript to all this, I'm not sure how long `diald' will be around. I keep hearing mention of the PPP package being able to handle demand dialing, and the latest Linux kernel quotes the PPP module as "version 2.3.7 (demand dialling)". At present I haven't looked into what this really means, or even if it is the same function performed by `diald'. Does anyone else know?

Again, if you have any comments or questions on what I've written, feel free to email on `frank@crawford.emu.id.au`. In particular, I am interested if there are ways I can improve this column, either format-wise or content.

❖

For the latest
news on
AUUG

Check out the AUUG
website at:

www.auug.org.au

News from the AUUG Business Manager

Liz Carroll
`busmgr@auug.org.au`

RETURNED TO SENDER...

Mail addressed to the following AUUG members has been returned. If any of you know their whereabouts, please contact AUUG on 1800 625 655 or email: `busmgr@auug.org.au`

Remember, if you are moving, please notify us of your new details - "Change of Address" forms can be found at the end of AUUGN.


We would really like to keep all our members up to date on the latest AUUG news!

Carbetta, Frank	Qantas Information Technology Airservices Australia
Desmond, Greg	TNT Australia
Donovan, Scott	NATSEM, University of Canberra
Erozan, Serdar	Australian Sports Commission
Formica, David	Computer Training & Consultancy
Fowler, Ian	SA Co-operative Bulk\ Handling
Freeman, Deanne	Philips Broadband Network
Gramp, Peter	CSIRO
Grunta, Franz	Workcover Corporation
Howard, Scott	University of Melbourne
Hoy, Craig	Australian Business Access
Hunt, Kevin	Joltech Pty Ltd
Kenyi, Tibor	Rio Tinto Technology Development
Leach, Bernard	
Lewis, John	
Liu, Xin	
McKay, Farrell	
Muirden, Richard	RMIT
Reeves, Scott	Hydrographic Sciences Pty Ltd
Saladino, Nic	Australian Electoral Commission
Shen, Bing	Queensland University of Technology
Stevens, Lorraine	Australian National Audit Office
Taylor, David	United Energy
Thuan, Le Dac	Vanoco Pty Ltd
Waterson, Howard	Mitsui Computer Limited
Whiting, Simon	Primary Industries and Energy
Wyss, Steven	Email Petroleum Systems

❖

Book Reviews

Sub-editor:
Mark Neely
mpn@infolution.com.au



CRACKING DES: SECRETS OF ENCRYPTION RESEARCH, WIRETAP POLITICS AND CHIP DESIGN

JAVA 2 UNLEASHED

Jami Jaworski
MacMillan Computer Publishing
ISBN: 0-672-31631-5

Reviewed by:
Peter D. Gray
pdg@dracl.its.uow.edu.au

It was with some dismay that I unpacked this book when it arrived. This book is big. Really big. It would be a useful item in hand-to-hand combat. There are 1424 printed pages. The table of contents alone occupies 25 pages. It looks like it should be comprehensive.

Perusing the table of contents tends to confirm that this book covers everything you would ever want to know about Java 2. It looks like a one stop reference. Unfortunately, this turns out to be not the case. Read on.

The introduction clearly states that this book is not for people wanting to learn java. It positions itself as a aid to programmers already familiar with the base language who want to expand their skills into new areas, such as RMI, swing or drag-and-drop.

The chapter layout is sensible, with the order well thought out. The layout of the book is easy to read, the typographical conventions easy to follow and the code fragments relatively easy to understand. The index is not as comprehensive as I would like.

The book covers just about everything even remotely connected with java, including topics of marginal interest to most programmers, such as the java speech and telephony API's. There are 6 appendices which include a language reference and an API description.

Listing out the topics covered would produce a long and probably boring list. Suffice to say that no matter what aspect of java you wish to inquire about, this book will have something on that topic.


How useful it will be is another matter.

The problem with this book I believe is simply that it tries to cover too much. The result is that nothing is covered in enough detail. It hardly ever delves into explaining underlying concepts, replying instead on giving an example of how you could perform a task using the tools being presented. I found this approach rather unsatisfying.

Perhaps this book is being too ambitious. Java may have reached the point where it requires a separate book on each of the major components.

If there is a niche for this book, it must be people who just want an overview of what java is capable of. In that respect it does demonstrate the almost mind-boggling explosion of java capabilities.

❖



BUILDING SCALABLE DATABASE APPLICATIONS OBJECT-ORIENTED DESIGN, ARCHITECTURES, AND IMPLEMENTATIONS

Peter M. Heinckens
Addison-Wesley, 1998
ISDN 0-201-31013-9

Reviewed by:
Michael Haldey, PhD
mhaldey@ozemail.com.au
and
Alex Juravlev
alexj@zip.com.au

The reviewed book, published as part of the Addison-Wesley Object Technology Series, deals with the design of complex database applications using a three-tier model.

The model is a natural evolution of the "classic" two-tier client/server model, featuring a middle layer responsible for the implementation of business rules. It attempts to address the problem of the close front/back end dependency that can often make design messy and complicate support, maintenance and further development of information systems.

The model presents a framework for the design of a persistence architecture that allows the separation of object-oriented business models and relational database models. Central to the architecture is the middle layer. It implements an object-oriented business model and communicates to the lower database layer through a set of

objects, allowing abstract database implementation and providing mapping of the object-oriented business model to the database model.

The book consists of four parts and an Appendix. Part One is an introduction to database community, the main models and issues of database application design. Traditional client/server two tier applications are also considered, along with a discussion of the problems frequently encountered when designing database applications using traditional methods. Object-oriented software design methods and issues are also discussed briefly.

In Part Two, a framework for a persistent architecture is built. The author considers building persistence into objects, abstracting database mechanisms and encapsulating data access.

Part Three deals with implementation of object business models, incorporating the persistent architecture previously suggested. This part of the book is fairly technical, with a lot of example code and a case study involving the implementation of a telephone directory. Special attention is paid to the inheritance of persistent objects and the use of associations. Also considered are the issues of transaction management and user interfaces.

Generally, the book is well written with a lot of examples and a few case studies, including a C++ implementation of a simple application, SCOOP, in the Appendix. The author gives a lot of practical advice and highlights typical design errors.

The book is not without its drawbacks. The C++ code examples are sometimes not well structured and used too frequently, which makes reading difficult for those not proficient with the language.

There are also a few confusing moments in the case studies. For instance, the persistence mechanisms are derived first from PSet class and later from IdObject class, without an attempt to clarify the change. The discussion of the complexities in the development of the middle layer, such as concurrency, is fairly shallow and virtually no attention is paid to such important issues like tuning, data caching, and others.

The author does not mention any commercial implementations of the persistence architecture. Though it was feasible some time ago to implement such framework from scratch for medium to large projects, nowadays the use of commercial implementations, such as the Java Beans persistence server, must be considered. Notwithstanding the above, the book will be useful for IT students, developers and designers of information systems. Parts of the text are a worthwhile reading for project managers and system architects.

❖

Victorian ACS News...

"IT TRENDS - TOWARDS THE NEW MILLENNIUM"

Intensive one day conference including overviews and case studies on E-Commerce, Data Warehousing, Knowledge Management and Virtual Private Networking.

Time:	9:00am - 5:00pm
Venue:	Le Meridien at Rialto, 495 Collins Street, Melbourne VIC 3000
Cost:	\$495 ACS and MIN members
Non Members:	\$595
To register	email: acsvic@acslink.net.au phone: (03) 9690-8000

Conference Update:

Open Source - AUUG'99

David Purdue
David.Purdue@auug.org.au

The program for "Open Source - AUUG'99", Australia's first conference dedicated to Open Source issues, has more or less been finalised. You should have received a registration form in the mail, and registration details (including program changes as they come up) can be found on the AUUG web site:

<http://www.auug.org.au/winter/auug99/>

Early bird registration will be available until the 20th of August, and the conference and tutorials are being held in the week of 6th - 10th September.

Highlights of the conference will be presentations by Eric Raymond, noted open source observer-participant anthropologist, Adrian Cockcroft, Sun Microsystems performance guru, Theo de Raadt, leader of the OpenBSD project, and Andrew Tridgell, the creator of Samba.

The AUUG AGM will be held at the conference, and the conference committee is open to suggestions for any BOF sessions you would like to see run - please send ideas to auug99@auug.org.au.

I ask that you, as AUUG members, please spread the news of this conference to as many people as you can who may be interested. A successful conference will ultimately give AUUG the resources to serve you, the membership, better.

Thanks,

David Purdue
President
AUUG Inc.

❖



Conference News:

What's on at USENIX

Full tutorial and technical session programs, and online registration, are available at <http://www.usenix.org/events/> for the following:

2ND CONFERENCE ON DOMAIN-SPECIFIC LANGUAGES

October 3-6, 1999, Austin, Texas

In cooperation with ACM SIGPLAN and SIGSOFT

Paper submissions due: March 22, 1999

2ND USENIX SYMPOSIUM ON INTERNET TECHNOLOGIES AND SYSTEMS

October 11-14, 1999, Boulder, Colorado

Co-sponsored by IEEE Computer Society

Technical Committee on the Internet

Extended abstracts due: March 25, 1999

Just added:

LISA '99--13TH SYSTEMS ADMINISTRATION CONFERENCE

November 7-12, 1999, Seattle, Washington

Co-sponsored by SAGE, the System Administrators Guild

Extended abstracts and Invited Talk Proposals due: May 25, 1999

TCL/2K: 7TH USENIX TCL/TK CONFERENCE

February 14-18, 2000, Austin, Texas

Extended abstracts due: September 1, 1999

❖ ❖ ❖

The USENIX Association supports user groups worldwide. We have recently inaugurated several programs to further this goal:

CO-SPONSORSHIP OF CONFERENCES & INTERNATIONAL AFFILIATE PROGRAM

There are many different models for affiliate membership or co-sponsorship; contact the Executive Director at ellie@usenix.org

INTERNATIONAL SPEAKERS PROGRAM

http://www.usenix.org/membership/intnl_speakers.html

USER GROUPS ON THE USENIX WEB SITE

<http://www.usenix.org/membership/ugs.html>

❖

Open Source - AUUG'99

8-10 September 1999
Carlton Crest Hotel, Melbourne

Full Tutorial and Conference Program along with Registration Forms
can be found at

<http://www.auug.org.au/winter/auug99/>

or call AUUG on
1-800-625 655



Representing the largest Unix and Open Systems event held in Australia and the inaugural Open Source conference, this event offers an unparalleled forum in which ideas and experiences will be presented to an audience with a major influence on the direction of computing in Australia.

Open Source - AUUG'99 will be preceded by two days of tutorials, and the conference from 8-10 September will see a mix of local and international speakers.

Open Source - AUUG'99 will address the following areas:

- Open Source projects
- Applications of Open Source Software
- Business cases for Open Source
- Technical aspects of Unix, Linux, and BSD variants
- Systems Administration
- Computer Security
- Performance Management and Measurement
- Networking, Internet (including the World Wide Web)
- Business Experience and Case Studies

Speakers include:

Eric Raymond presenting "Anthropology of Open Source"

Eric is an observer-participant anthropologist in the Internet hacker culture. His research has helped explain the decentralised open-source model of software development that has proven so effective in the evolution of the Internet. His own software projects include one of the Internet's most widely-used email transport programs.

Theo de Raadt presenting "strlcpy and strlcat - Consistent, Safe, String Copy and Concatenation"

Theo is the leader of the OpenBSD project. OpenBSD is an open source operating system (based on BSD 4.4lite and NetBSD) that concentrates on aspects of security. Theo will talk about the OpenBSD approach to managing an Open Source project.

Jon "maddog" Hall presenting "Linux International"

Jon is the executive director of Linux International, dedicated to promoting the use of Linux. Working with computers for over 25 years, 19 with UNIX, Jon held various positions with Digital, prior to which he was a Senior Systems Administrator in Bell Laboratories' UNIX Group. Jon has also been directly responsible for the port of Linux to the Alpha processor.

Adrian Cockcroft presenting "High End UNIX Directions"

Adrian is Sun Microsystems' performance guru. It is said he knows more about how to tweak that last 1% performance out of a system than anyone on the planet. He will be talking about the state of play with performance tools on UNIX, and on aspects of server consolidation and resource management. He will also present his performance tuning tutorial.

Andrew Tridgell presenting "rsync in http"

Andrew is best known as the author of SAMBA - an open source tool that allows UNIX systems to talk the network language of Microsoft Windows machines, namely the SMB protocol. Andrew now leads the Samba Team, who carry on development of Samba.

He will be speaking on how to manage a large scale open source project, and will also be presenting a tutorial on the internals of Samba and the SMB protocol.



Sponsorship Opportunities:

Open Source AUUG'99

*AUUG'99 Conference
September 8-10, 1999
Carlton Crest Hotel
Queens Road
Melbourne VIC Australia*

THEME: "OPEN SOURCE"

The 1999 AUUG winter conference will be held at the Carlton Crest Hotel, Queens Road, Melbourne, Australia, between September 8th and 10th.

The conference will be preceded by two days of tutorials, on September 6th and 7th.

Information on papers can be found at the AUUG website :

www.auug.org.au

There are also many advertising and sponsorship opportunities available for your Organisations. Information is available from the AUUG business manager.

To take advantage of these, please contact:

Liz Carroll
AUUG Business Manager
Phone: 02-9858-4542
Email: busmgr@auug.org.au



OPEN SOURCE AUUG '99

Carlton Crest Hotel
Melbourne
8-10 September 1999



DIAMOND SPONSORSHIP

A\$10,000

Includes:

- 2 complimentary registrations for the Conference
- 2 complimentary invitations for the Cocktail Reception
- 2 complimentary invitations for the Conference Dinner
- logo displayed in conference Plenary Hall
- acknowledged on all appropriate occasions in both print and verbally
- small display area

Listed and identified as a sponsor in:

- the conference brochure
- the conference final program with company description

CHOICE OF:

- **Conference Brochure**
 - wide distribution to key decision makers
 - areas of exclusive advertising
 - immediate impact prior to the conference
- **Welcome Reception**
 - prestigious event allowing sponsor to make first impression on the delegates
 - reception identified as being sponsored by the XYZ company on all printed material
 - signage on the evening
 - opportunity to address delegates
- **Conference Dinner**
 - dinner identified as being sponsored by the XYZ company
 - name printed on dinner menu
 - opportunity to distribute mementos and address to the audience
 - banner identifying the sponsoring company
 - name of sponsoring company on entry tickets



PLATINUM SPONSORSHIP

A\$7,500

Includes:

- 1 complimentary registration for the conference
- 2 complimentary invitations for the cocktail reception
- 2 complimentary invitations for the conference dinner
- logo displayed in conference plenary hall
- acknowledged on all appropriate occasions in both print and verbally
- display space

Listed and identified as a sponsor in:

- the conference brochure
- the conference final program

CHOICE OF:

- **Conference Proceedings**
 - 2 A4 pages of exclusive advertising
 - long term usage and shelf life as it is a reference material
- **Tee-Shirts**
 - offering long term usage and company message to recipient
- **Conference Satchel**
 - Satchel offering long term usage and company message to recipient



GOLD SPONSORSHIP

A\$5,000

Includes:

- 1 complimentary invitation for the cocktail reception
- 1 complimentary invitation for the conference dinner
- logo displayed in conference plenary hall
- acknowledged on all appropriate occasions in both print and verbally
- display space available for one day at the conference

Listed and identified as a sponsor in:

- the conference brochure
- the conference final program

CHOICE OF:

- **Keynote Sessions**
 - opportunity to introduce the keynote session
- **Lapel Badges**
- **Speakers Reception**
 - event allowing sponsor to make first impression with speakers
 - reception identified as being sponsored by the XYZ company on all printed material
 - signage on the evening
 - opportunity to address speakers



SILVER SPONSORSHIP

A\$2,500

Includes:

- 1 complimentary invitation for the cocktail reception
- logo displayed in conference plenary hall
- acknowledged on all appropriate occasions in both print and verbally
- rack space for promotional material

Listed and identified as a sponsor in:

- the conference brochure
- the conference final program

CHOICE OF:

- **Advertisement in conference program**
 - reaching an audience of key decision makers
- **Conference folder insert**
 - individual inserts in conference satchels
- **Advertisement in conference proceedings**
 - A4 size



ADDITIONAL OPPORTUNITIES

- Audio Visual
- Conference Network

Further information is available from the AUUG Business Manager, Liz Carroll.

Contact details:

email: busmgr@auug.org.au
Phone: 02-9858-4542



ADVERTISING

Various advertising opportunities are available from \$1,500, including inserts on the brochure racks, handing out promotional material etc.

Further information is available from the AUUG Business Manager, Liz Carroll.

Contact details:

email: busmgr@auug.org.au
Phone: 02-9858-4542

TELLURIAN

Tellurian Pty Ltd

Come to us if you need seriously capable people to help with your computer systems. We're very good at what we do.


- Unix, Macintosh and Windows experts
- Legacy system re-engineering and integration
- System management and support
- Internet access

Our two current major projects:

- Support and development of an integrated environment covering applications running on IBM3090, DEC Alpha, SCO Unix and Nortel switches. Just imagine the cost benefits of supporting over 500 concurrent users on four little 486 and Pentium PCs.
- From the ground-up implementation of MFC and Windows API on Apple Macintosh. We've got our client's Windows MFC application running, bug-for-bug, on Apple Macintosh.

Tellurian Pty Ltd
272 Prospect Road
Prospect SA 5082

(08) 8408 9600
www.tellurian.com.au
sales@tellurian.com.au



***10% DISCOUNT
TO AUUG MEMBERS
ON OUR COMPLETE RANGE
OF COMPUTER, BUSINESS
AND GENERAL BOOKS***

**AMERICAN
BOOK STORE**

173 Elizabeth St, Brisbane Queensland 4000
Ph: (07) 3229 4677 Fax: (07) 3221 2171 Qld Country Freecall: 1800 177 395
american_bookstore@compuserve.com

Name: _____ Date: _____

Address: _____
 _____ Post Code: _____

Phone Number: _____

Payment Method: Cheque Money Order Amex Bankcard
 Diners Mastercard Visa

Card Number: _____

Expiry Date: _____ Signature: _____

This is a: Special Order Mail Order Book on Hold

QUANTITY	TITLE	PRICE

SUBTOTAL \$
 LESS 10% DISCOUNT \$
 POST & PACK \$
 TOTAL \$

POSTAGE AND HANDLING FEES: 1 BOOK \$6.00 2-4 BOOKS \$7.00
 BOOKS OVER \$70.00 WE WILL SEND CERTIFIED - PLEASE ADD ANOTHER \$1.50 OR WAIVE
 CERTIFIED DELIVERY.

FOR SPECIAL ORDERS, PLEASE ENCLOSE \$10.00 PER BOOK AS A DEPOSIT.

From the pages of unigram·X

WITH SOURCEXCHANGE, FREE SOFTWARE IS NO LONGER UNPAID LABOR

By Rachel Chalmers

The way companies develop software could be about to change. O'Reilly & Associates and Hewlett Packard Co have introduced SourceXchange, a web-based service that links corporate sponsors with open source development projects. The companies plan to launch the service in July with seed projects funded by HP. If all goes well with these betas, other enterprise sponsors will be encouraged to sign up. "HP came to us with a challenge," explained Brian Behlendorf, O'Reilly & Associates chief technology officer and one of the founders of the open source web server, Apache. "They wanted to further the development of specific open source software projects, but they needed a better way to reach out to the wider developer community and coordinate the development process."

SourceXchange was put together under Behlendorf's direction. It maintains a database of all published project Requests for Proposals (RFPs) posted by corporate sponsors. Developers and their teams are registered with the site, which manages RFP responses from the community and arranges payment for developers. SourceXchange will also incorporate peer review and project milestones to help ensure quality and reliability. For the sponsors, SourceXchange could provide the outsourcing option they need to extend software development resources and budgets. For developers the benefit is even more clear. SourceXchange provides a change to work on open source projects for cash!

The process begins when a sponsor needs some open source code written for a particular purpose: a device driver, a Perl utility or an Apache module would be ideal. The sponsor registers with SourceXchange and agrees to shoulder development costs. Now the sponsor needs to create an RFP and submit it to the SourceXchange web site. Internal peer reviewers will look over the RFP and work with the sponsor to refine it, if necessary. Next, the SourceXchange administrator posts the RFP to the site. A set comment period follows before the RFP is closed.

Finally, developers who have registered at no cost to themselves will be invited to become active developers and respond to the RFP. A lead developer may work alone or recruit a team from the registered developer pool. The sponsor and a peer reviewer will look over the proposals and the sponsor will select one to accept. Development gets under way. The sponsor and peer reviewer keep track of milestones; the lead developer is paid as each milestone is achieved and approved.

When the project is complete, SourceXchange elicits feedback from all participants. The organizers promise to use this feedback to improve the process. Source code will be published under an appropriate open source license and the whole process will be archived for future reference, helping minimize the need to reinvent the wheel.

"All too often a great idea will come up while developing Apache and disappear because there is no one with the time to work on it," Behlendorf says. "Likewise, I've seen lots of companies asking for new features or ideas, willing to pay for the work, but without the personnel to help make it happen. So there's this tremendous 'potential energy' building up, keeping a lot of useful code from being written, and many of us thought this was a problem worth fixing." If SourceXchange can indeed unite the open source development community with the resources of enterprise software vendors, the problem will be well and truly solved.



LINUX VERSIONS OF NOVELL NDS, BEA TUXEDO AND WEBLOGIC

By the end of calendar year 1999, Novell Inc's Novell Directory Services (NDS) will run natively on Linux. The directory technology should make it easier to integrate Linux resources with those residing on NetWare, NT and Solaris systems. Customers will get a single infrastructure for managing and securing their platforms and applications. NDS supports single sign-on and single point of administration, which helps reduce network maintenance costs. Novell's announcement adds fresh credibility to Linux as an enterprise platform, as does a report in Infoworld that BEA Systems Inc is poised to announce versions of its Tuxedo distributed transaction processing software and WebLogic application server for Red Hat Linux 5.2.



STALLMAN SOUNDS OFF ON SOFTWARE PATENTS IN EUROPE

Richard Stallman, founder of the GNU project, has spoken out against European Patent Office plans to extend European patent law to software. The EU wants to draft pan-European legislation to make intellectual property protection consistent across member nations, but Stallman insists that this would be disastrous for the free software community and for developers in general.

"There is a push now to make Europe have software patents like the US. If this happens, it could be disastrous for free software (and bad for anyone who wants to develop software in Europe)," he told the Freepatents.org web site. He called for developers and other concerned parties to write to European officials: "If two countries object to the plan, that will be enough to block the plan."



LINUXCARE SNAPS UP ANDREW TRIDGELL, DEVELOPER OF SAMBA

Linuxcare Inc has named Andrew Tridgell, the Australian academic and team leader for Samba, to the position of senior researcher. Samba provides a complete replacement for the networking services included in Windows NT, OS/2 Warp, Sun NFS and Novell NetWare servers, letting Linux perform in their place. "Samba is a virtually ubiquitous open source application, depended upon by almost every Linuxcare customer," explained Linuxcare chief technology officer David Sifry. All this makes Tridgell is something of a catch. He will continue to serve as Samba team leader and will remain a visiting fellow at the Australian National University.



SCO TAKES SYSTEM V SYSTEM MONITORING OPEN SOURCE

Santa Cruz Operation Inc is turning over another piece of Unix System V code to the open source community, the sar system monitoring suite. Under the project, called Osar, source code will be released as soon as proprietary or non-portable code is removed. SCO wonders whether Sun Microsystems Inc will be so keen to promote sar as it has the SCO's lxrun Linux-on-Unix runtime libraries.



SIEMENS MAINFRAME USERS SHOWN OPEN FUTURE

Siemens Computer Systems says its BS2000/OSD has become the first mainframe OS to receive "internet server" branding from The Open Group. A company spokesperson said that the open branding should give customers the confidence to migrate from Siemens' Fujitsu-supplied CMOS mainframe platforms, to its own MIPS-based RISC platforms and to Intel IA-64 platforms once they become available later this year. The move should ensure some extra shelf life for BS2000 environments, which are mostly limited to the operations of a few hundred large German customers, but Siemens is also hopeful that the MIPS and IA-64 escape routes might even encourage a little growth among the 3,000 BS2000 licenses currently in the market.



SUSE RELEASES LINUX 6.1 FOR ALPHAS

German commercial Linux distributor SuSE has started shipping the Alpha processor version of its Linux 6.1 bundle. The Alpha version includes YaST (Yet another Setup Tool), which integrates maintenance and configuration tasks such as

updating kernels, performing backups and administering users and groups. The Alpha implementation of Linux 6.1 also includes SaX (SuSE advanced X configuration), which configures most of the I/O devices on an Alpha workstation or server. SuSE says that the Alpha version, like its X86 version, comes with a manual and 60 days of support. In the future, SuSE says it will announce X86 and Alpha versions of its Linux distributions at the same time.



AMD PRODUCES FIRST COPPER K6s, 1GHZ ATHLON TO FOLLOW

Advanced Micro Devices Inc has produced its first copper K6 chips and is expecting to produce 1GHz Athlon CPUs using copper interconnects by the end of the year. The chips are being produced by AMD's new production facility, 'Fab 30', based in Dresden, Germany.

In interview with the German web site Heise Online, the plant's manager Hans Deppe, said that he was happy with the first K6 batches using the copper technology but would not reveal what the yield was. Staff at the plant expect the first batches of Athlons with clock speeds of 1GHz or over by the end of this year or early 2000. The Dresden plant was opened in February of this year. It has a staff of 850.



TURBOLINUX J 4.0 OUTSELLS WINDOWS 98 AND MAC OS IN JAPAN

TurboLinux Inc, formerly Pacific HiTech Inc and the Linux distribution powerhouse of the Asia-Pacific region, is boasting that its newly released TurboLinux J 4.0 Workstation outsold Windows 98 and all other Linux operating systems in retail sales in Japan. Market research company Business Computer News analyzed sales at 200 computer stores throughout Japan and declared TurboLinux J 4.0 the number one seller. In its first week of sales, TurboLinux achieved a 20.09% market share, ahead of Windows 98 Upgrade on 13.25%, Mac OS on 10.23%, Windows 98 on 9.15% and all other competitors with less than 7% apiece.



ALPHA READIES LOW-COST MOTHERBOARD

Alpha Processor Inc (API) is planning to offer a low-cost Alpha chipset in the fall. The UP1000 is based on Advanced Micro Devices Inc's Iron Gate chipset. It shares a number of AMD Athlon-compatible features such as a 'slot B' connection and a 21264 Alpha bus. The motherboard has 2Mb of Level 2 cache memory. There are no prices as yet but it should be cheaper than its \$6,600 sibling, the UP2000. The board is expected to be available in September or October.



LINUXPC IS NEXT ON THE AGENDA FOR VA RESEARCH

By William Fellows

VA Research Inc is determined not to let Caldera Inc, and especially LPO (linux public offering) frontrunner Red Hat Inc, cast a shadow on it and it has been on the road taking soundings about what we'll refer to as the LinuxPC, as well as touting its value-add, hardware systems.

It's going to be tough though. In addition to the IPO, Red Hat last week garnered more funding, this time from Novell Inc's internet fund, effectively becoming anointed Novell's Linux directory partner (see separate story). The company has also rolled up an e-commerce product combining Red Hat Linux 6.0, Apache web server, RSA Data Security 128-bit, Netscape roaming module, Squid proxy and cache server, Webalyzer server log and a \$25 discount to use with Thawte digital certification. It also has a separate e-commerce directory within its Linux distribution with trial versions of HP, CCVS and MiniVend e-commerce software. Its idea is to capture businesses and ISPs who want to focus on sales rather than their web site, Red Hat says. It's priced from \$150.

There's a mixed bag of news on other Linux fronts. The Sun- Netscape alliance (iPlanet) has supposedly nixed putting the forthcoming Netscape Application Server 4.0 up on Linux; after 4.0 it gets souped up with Sun's NetDynamics. From the beginning the alliance had said that Linux was to be an important platform for it. VA Research CEO Larry Augustin claims Sun is getting hurt by Linux at the low-end. The handful of stalwarts left at Mozilla surely can't pick that one up. However, Augustin claims that the company is getting so many Linux boxes out of the door that it didn't actually need the \$25m venture funding it raised last month. It's effectively running on cash. But it will wait to see how the Red Hat IPO goes off - slated for August 10 - until deciding whether to cast its stone to the market.

The 150-person company claims it is doubling sales every quarter and that this month it will have shipped more systems than the entire calendar second quarter. Most sales are one- and two- processor systems to dotcom companies and ISPs. Augustin says the biggest opportunity for Linux is going be IA-64 because of the huge increase in memory and caching it will provide. That's what these companies are crying out for, he claims. It show its product running on a Merced simulator at August's Linux World jamboree.

VA Research has investment from both Intel Corp and Silicon Graphics Inc but can't say how or what it has contributed to SGI's forthcoming Linux servers even though it admits it has been working with SGI on the products. Moreover, it doesn't see SGI's Linux roll-out as a threat. "Let's wait until we see it," said Augustin.

While getting the enterprise functions on to Linux is important, including a journaled file system, failover, high availability, and clustering (for Oracle Parallel Server), that is pretty much a done deal.

The big thing in Linux is the desktop, VA insists. It's what everyone in the community is working towards. It needs ISVs to buy in to really get the opportunity going. What that opportunity would mean in terms of products also remains to be seen. Augustin believes a simple LinuxPC with business/office applications is the grail. Moreover, if the Red Hat and other LPOs go off strong then Linux is likely to strike a chord with US consumers, currently obsessed with stocks, the web and technology.

Linux is a no-brainer for the CAD and workstation ISV community, Augustin believes. Desktop Linux needs a clean user interface before it has a chance and several vendors already have designs. VA says it may show off its work in the area at Linux World, and possibly its implementation of X-Windows too. Augustin also thinks the embedded space is an opportunity given that Unix missed the boat there. The opportunity couldn't be better put by the arch enemy itself: "The Unix phenomenon is scary. It doesn't go away. Linux is a serious, albeit crazy, implementation of Unix on the Intel platform," said Microsoft president Steve Ballmer.



LINUX DOESN'T MAKE THE GRADE ON SIEMENS' EIGHT-WAYS

Siemens Computer Systems is gearing up for the introduction of its eight-way Intel Xeon Primergy servers (no they still haven't gone GA yet) but they'll be without Linux, which the company had hoped to supply. They will run NT and SCO UnixWare. Sources say Linux benchmarks on Primergy were so catastrophic that Siemens cancelled the plans it had to support the operating system, at least until there is a robust multiprocessing version of Linux available. Siemens works with Red Hat and SuSE in Germany. It's not a Primergy issue. It offers Linux on quads.



COMPAQ SHOWS OFF DS20 CLUSTER, NEW SWITCH FABRIC

By Timothy Prickett Morgan

Compaq's Digital unit made some waves with a parallel supercomputer cluster made up of 128 relatively cheap two-way DS20 servers. The machine, which was built under a four year, \$11 million contract from the US Department of Energy for its Accelerated Strategic Computing Initiative, or ASCI, is the fastest supercomputer ever put together by Compaq/Digital and ranks as number 49 on the Top 500 supercomputer sites in

the world. The goal of the ASCI program is to develop supercomputers that use off-the-shelf technologies that have 30Tflops of power by 2001 and 100Tflops of power by 2004. While the DS20 cluster fell far short of those goals with only 154.4Gflops of computational power (that really is quite a lot), the cluster demonstrated that Compaq can build big enough machines to start building a business in the high performance computing market that goes beyond workstations, SMP servers and normal TruClusters. The DS20 cluster was five times as fast as the AlphaServer 8400 clusters, each with 84 processors, installed at Sandia National Labs and Lawrence Livermore National Labs, which are the fastest machines Compaq has ever created. The cluster is considerably more powerful than the installed Sun Starfire or HP V2500 minisupers installed around the world -- although Sun and HP can demonstrate similar number crunching power if they want to.

The fact is, Compaq has a decisive advantage with the speed of the Alpha 21264 processors, which are the fastest math majors in the chip business, but it needs to prove to potential customers who up until now might go with an IBM RS/6000 SP, Hewlett-Packard V2500 or Sun Starfire for a baby supercomputer that Compaq has the know-how and the application software to make such supercomputer clusters practical. It is one thing to show the DOE that you can build a machine that supernerds can program on. It is quite another to build a commercial supercomputer business, as IBM has done brilliantly with the RS/6000 SPs over the past six years and as HP and Sun are starting to do with their respective V2500 and Starfire servers. Compaq says that because of the work it has done under the DOE's ASCI project, it will be able to bring HPC technologies to market two years earlier than it had originally planned.

One of the key technologies in Compaq's TruCluster bag of tricks is the Memory Channel interconnect, but for this breed of clusters, Compaq is working with a new partner, Quadrics Supercomputing World, Ltd, an HPC clustering specialist that hails from Bristol, England, to build the interconnect fabric. With the first generation of TruCluster's Memory Channel hardware and software, message passing between nodes had a latency as low as 8.5 microseconds. The channel offered bandwidth of up to 60Mb per second. With the forthcoming Memory Channel 2 interconnect, Compaq will be able to show 100Mb per second full duplex message links between nodes and 2.1 microsecond latencies. But the MC2 interconnect, like its predecessor, will still be limited to a maximum of eight nodes. This isn't good enough for the supercomputer market, especially when Compaq won't have the 16-way "Wildfire" AlphaServers using the 767MHz EV7 Alpha processors ready until the fourth quarter, and it won't have 32-way or 64-way follow-ons to the Wildfires ready until next year or later. Compaq can't count on SMP to build big servers any more than can IBM, Sun or HP, and that is why it needs better clustering. That is why Compaq has cut a deal with Quadrics.

According to Quadrics, which expects Compaq to start delivering production Alpha clusters using its Qnet switches starting sometime in the third quarter, its fabric will allow more scalability than and similarly low latencies to the MC2 fabric. On the DS20 cluster using the Qnet fabric, which consists of a network interface card that sports a custom ASIC designed by Quadric in each machine that is connected to a central 16-port switch, the cluster was able to show latencies on the Linpack benchmark test as low as 5.5 microseconds and 200Mb per second of message passing bandwidth. Quadrics says that its interconnect can achieve 340Mb per second bandwidth, although it runs up against PCI bandwidth constraints and hence generally tops out at around 200Mb per second, just as it did in the DS20 cluster Compaq just tested. (That will change when Compaq soon ships an improved PCI adapter). Quadrics says that its Qnet fabric can provide latencies in the 2.5 to 5 microsecond range, which is comparable to Compaq's MC2 interconnect. The big difference is that Qnet supports 128 nodes, and its architecture is extendable to 2,000 nodes or higher.

If Compaq can build a 128 node cluster with DS20s, it won't be long before it builds one with ES40 four-way motherboards and thereby doubles the power of its fastest supercomputer, and it won't take long to push it to 256 nodes and quadruple the power. That would put Compaq solidly in the Top 15 supercomputer sites -- if anybody bought one -- and nipping at IBM's heels with its ASCI Blue RS/6000 supercomputer at Lawrence Livermore Labs. Even still, a 1,024 processor Alpha cluster would still be a far cry from the 2.1Tflop ASCI Red super at Sandia Labs, which was built by Intel, or the 1.6Tflop ASCI Blue Mountain machine at Los Alamos, which was built by SGI. Compaq would have to build a 1,000 node, 4,000 processor machine using today's Alpha 21264 chips to meet or just barely beat ASCI Red.


Customers who want more modest supercomputers from Compaq won't have to wait. Last month, Compaq announced two baby supercomputers based on its E40 four-way Alpha boards called the HPC160 and the HPC320. Both machines use the 500MHz Alpha 21264 processors. The HPC160 is a 16-way server with four ES40s linked together with a Memory Channel crossbar switch and 8Gb of memory. The HPC320 is a 32-way server with eight ES40s, 16Gb of memory and two crossbar switches lashing together. Both machines offer up to 128Gb of memory. Compaq hasn't provided benchmark results on either machines as yet, but a fully loaded HPC160 should be in the 10Gflops range.

❖

New

Addison
Wesley
Longman

from Addison Wesley Longman



Mastering Java 1.2


John Zukowski

Mastering Java 1.2 provides the most in-depth coverage of JDK 1.2, enabling the reader to take immediate advantage of the many new and improved features of the latest standard from SUN.

Inside this book there's ample discussion of the basics—with everything that makes Java ideal for simple desktop and Internet applications. Experienced Java programmers will benefit most from the detailed coverage of what's new, including the many capabilities that make Java 1.2 ready for true enterprise development.

ISBN: 0-7821-2180-2

\$99.95 w/CD-Rom



Linux Application Development

Michael K. Johnson and Erik W. Troan

This practical reference guides programmers developing Linux applications or porting applications from other platforms. Linux is fundamentally similar to Unix—therefore, much of the book covers ground familiar to Unix programmers—but this book consistently addresses topics from a Linux point of view. The aim throughout is to provide the detailed information you need to take full advantage of Linux.

ISBN : 0-201-30821-5

\$99.95

AUUG Members receive a 20% Discount off the RRP

order form

To order a copy of these texts, please fill in the details below.

This form should be faxed or mailed (freepost) to Cassy Nacard, Trade Marketing Coordinator at Addison Wesley Longman, Level 1, 2 Lincoln Street, Lane Cove NSW 2066. Ph: (02) 9428 8086 or Fax: (02) 94279922

Book Details

Author/Title	ISBN	Price
_____	_____	_____
_____	_____	_____
_____	_____	_____
	Total less 20%	_____

Order details

If you are **purchasing** the books detailed above, please give the following information. (Please note that books can also be ordered directly from bookstores.)

Book(s) price: _____ Note: Prices are subject to change without notice

Postage: 1 book = \$4.00, 2 or more = \$7.00 Total \$ _____

I enclose a cheque made out to Addison Wesley Longman

Please charge my credit card:

Bankcard Visa Mastercard American Express ID No: _____

Card Number: _____ Expiry Date: _____

Signature: _____

Your Name: Mr/Mrs/Ms _____

Address: _____

Daytime Ph: _____

Minutes of:
AUUG Exec Meeting

DRAFT FOR DISCUSSION

Date: 30/7/99
Time: 10:00 AM to 4.00 PM
Location: Aurema
79 Myrtle Street
Chippendale

Attendees:

Lucy Chubb	LC
Luigi Cantoni	Lca
Malcolm Caldwell	MC
Elizabeth Carroll	EC
Gunther Feuereisen	GF
David Newall	DN
David Purdue	DP
Michael Paddon	MP

1. Apologies

Mark White	MW
Stephen Boucher	SB
Peter Gray	PG

2. New Executive Committee

Welcome to David Newell by DP.

DP takes over from LC
Lca takes over from SB
SB takes over from MW

Will discuss Exec Committee Duties later.

All Exec to supply new contact details to DP.
Action All

3. Presidents Report

First stint at president, jointed the ranks of Greg Rose, now held all the roles, with the exception of Treasurer.

Does not regard himself as a natural leader, but a good organiser and good at meetings. Would like to run exec meetings - free flowing, but keep moving and keep to agenda. Outside, the President should be the face of AUUG, and would like to get AUUG's opinion out to where it is noticed. Views about what a user group should be for; and the fact that we are a society under the terms of the act puts a responsibility on both us and the members.

Plans for the year - create more of a sense of community among members, provides services via electronic services etc.

Big thank you to Lucy as President - standing in front of the fan from which the shit, etc, etc.....!!

Had a lot of responsibility due to various issues. We are now well into a better way to process things - and forming better relations with members and suppliers.

Move a vote of thanks to Lucy.
Moved GF/MC

Motion to accept:

Accepted: LCa/GF

4. Secretary's Report

Both incoming and outgoing secretaries are absent.

5. Treasurer's Report

Outgoing Treasurer is absent. - New Treasurer will fly to Melbourne on 2 August to do the handover.

Motion that AUUG pays for Lca to fly to Melbourne to do the handover.

Accepted: LC/GF

SB unsure of the account balances. Have a tax cheque coming up. Liz's entitlements, then Executive Expenses.

Lca wants to achieve a smoother flow of funds, a better methodology - as we are scattered around Australia. Wants to make sure the Chapters are happy - so it can work. Wants to work the conference so that it succeeds financially - wants to walk into it knowing what we will make.

Motion to accept:

Accepted: MC/LC

6. Business Manager's Report

Has moved the AUUG Secretariat office to her office. Currently working through the renewals, using Foxpro database - this may not be the best database of choice, however may look at purchasing Filemaker Pro which may work better.

Processed and banked approximately 35 renewals since taking over, however, there are approx 100 cards yet to be ordered for the June renewals which she will arrange to do. As these are late, a notice will be included in AUUGN.

Following the move, there are boxes to be put into storage - Lca will collect these.

As EC will be moving, AUUG will keep the current phone numbers for the next month, after which the Paddington phone number will have a recording giving the new office number, the 1-800 number will be forwarded to the new office, along with either a new office number, or if possible she will keep the existing 9858-4542 number.

AUUG will keep the PO Box at Kensington and arrange to have mail from it redirected to the Business Manager's address.

Other topics which the Business Manager will enlarge on and which she has been working on include, AUUGN, Chapter Activities, AUUG Survey, AUUG99 and AUUG 2000.

7. Secretariat's Report

Duties currently being handled by the Business Manager - currently processing the June 1999 renewals.

8. Minutes previous meeting

Accepted: LCa/MC

Action Items

All action items were covered, a separate document covers the status of these.

9. Events

AOSS

40 attendees. Reconciliation to be done - SB has the cheques, need to be banked. Came out in profit - good press, everyone said good things about it. Should do another one - do another outside of Melbourne - look at doing something a little more high profile - talk to Usenix. Need to look at the next city - MP would like to have a local involved.

Want to put up a web site: AOSS.auug.org.au
Action MP

Would like to put it up electronically.
MP will chase up the evaluation forms.

Java Symposium

Proposed a small scale roadshow by Jan Newmarch on Java 1.2 - EC hasn't received a response yet. As he is now in Brisbane, possible do Bris/Syd/Can.

MC has stated security should be run each year - in Sydney around November would be good. Alan Cowie is interested.

Contact Alan, then CFP - try to include it in AUUGN. Get date and location & cost. Register interest in speaking and attending. - EC talk to Alan. Do a half page ad. Decision to be made on Friday.

Look at an Open Source event early next year. David Newell to work with MP - try for mid February 2000 in Adelaide. DP to approach Usenix for funding. As well Linux Int. - get CFP in AUUGN.

QLD have their monthly events - suggestion for a May event?

Roadshow in March - aim for Jan Newmarch, coincide with Summer Events. Start working on this for yes from Jan this week, so we can include it in AUUGN.

10. Constitutional Changes

Proposed wording sent to Chris Maltby.

11. AUUG99

Most important challenge at this point in time. Needs good attendance and profit - if not successful, AUUG is in trouble. Have a reasonable program - in line with topic of interest to this community. Good drawcards - Adrian Cockcroft. Everyone needs to let everyone know its on. Luigi to tell the KL people. ACS advertising. Mark sending to SE Asia group. David Newell to contact IEEE and ACM.
Action DN.

DP will email Ellie.

EC talk to SMH and Age. Installathon should also be in the diaries as a separate event.

Also need volunteers for this.

Sponsorship

Chase up on 13 August. Talk in the meantime.

EC to chase sponsors in Melbourne

Try IBM - sponsorship - ad on proceedings
Action EC

Speakers and Program

Complete, only question is Jon Hall - to be confirmed by MW

MW would like to do a Keynote - State of the Nation Address

Would like to have SCO involved - DP would like to get him into the program.

Oracle are a question mark - therefore we may be able to use this spot for SCO.

Papers are coming in - however gave the wrong date for papers to be received. There will be a message out this weekend.

EC to design the cover, quantity - 300.

Budget

Domestic fares through the 35% discount.

Try for international on credit cards and we reimburse them.

PR & Press

Official publication will be out within the next 2 weeks

Charles Wright on 30L runs a program on computer problems - action MP to call in - be on the program to talk about Linux

Network

MP will talk to his management

DP will talk to Primus again - possible sponsorship

Other

Session Chairs - DP looking for volunteers - will email the chapters for this

12. AUUG 2000

Tricky year due to the Olympics.

Looking at July - we need to decide at this meeting.

Call for Papers to go in this edition of AUUGN

Discussed merits of Sydney, Adelaide, Gold Coast and Hobart

Gold Coast - anticipate 75 locals

Hobart - won't get 100 locals

Sydney - 75 ?

Adelaide - 75 x

Gold Coast International will be the venue - Mark White to recommend a Conference Chair, and AUUG to approach David Hughes for Program Chair.

2nd or 3rd week in July.

Theme for the Conference suggestions:

Open Source hits/on the road

Open Security voted the theme

13. PR & Press

Need some more writers for the press.

Bottle of wine up to \$10, and thank you letter from the President for writers -

Options

Authorise \$300 to pay for gifts for people who write for AUUG.

Accepted: GF/DN

14. Tasks and Responsibilities

Come up with a list of responsibilities for each committee member, go through this at the next meeting.

Roles include:

Member Champion
AUUGN - GF
Webmaster
Assistant Treasurer
Event Co-Ordinator
Lions Award and Student Program - LC

15. AUUGN

Hasn't received anything - EC to provide some info next week.

GSM article coming in from Tasmania.

Frank Crawford will be sending in an article, Michael Paddon will give Blackened Voodoo 2.

Wants to finish by the second week of August, it arrives at least a fortnight before the conference.

Traps and Tricks there is nothing.

DP will put in the President's column.

EC ask SMH if we can reprint articles we have submitted.

Hoping for last minute submissions over the weekend.

Chasing book reviews.

Minutes to be included

Result of member survey to be printed

Election results to be printed in AUUGN

Notice of AGM to be in

Missing members to be supplied by EC

Meeting summary as an agenda item for next meeting

16. Member Survey

Raw results sent by GF, will be reposted with corresponding questions

Comments included, Linux, accused of being a copy of SAGE, don't know what's happening in AUUG, AUUGN is thin, a need to know what the chapters are doing, Traps and Tricks a good draw card - need to receive some more.

Draw for the free conference registration result:

17. IDG

IDG are ready to go, just need AUUG to advertise it, needs to be in AUUGN and on auug-announce. Include details of how to subscribe. Action GF to put into AUUGN - GF will write it up, DP will send to auug-announce.

18. NT Chapter

Tried to hold a conference, over the past year they have lost their secretary, the treasurer became uncontactable. MC tried to arrange the conference, with the help of a couple of people, originally delayed, it was rescheduled. One week prior there were no registrations - 1 day workshop with 6 attendees at \$100. Two interstate speakers were on for free, the rest of the papers to be used for later meetings - attendees that showed up approx 12.

MC needs some re-imbursing from running the conference.

19. SA Chapter

MC has found two volunteers to start a chapter in SA. Greg Lehey is keen to start - EC to extract of SA members and provide to Greg. DN to contact him, let DN and MC know once he has received the member list.

AUUG will pay expenses for a meeting in Adelaide for forming a chapter. Accepted GF/LCa Once we have the requisite number of signatures, we can then have chapter formed.

20. Other Business

Tag Line - to be included as Agenda Item at next meeting.

Web Server - MW volunteered some hardware, MP bandwidth, next Exec meeting, appoint a webmaster.

Membership Cards - must start to put random code on them.

Open Source Accreditation, MP would like to aim at what makes AUUG relevant to people. Has always stood for openness - we now have opportunity to articulate this - seems that AUUG can help people who are trying to build open source or open systems software. Possibly define what is open source. Show them what is really open source. AUUG could possibly lead an industry wide debate on this and look at an Open Source stamp on it. Currently we do not do a lot that is visible in the industry. We do not believe there is a particular Australian open source group - AUUG appears to be the closest to it. Each time we come together we redefine ourselves, something we need to do to give benefit to our members. MP to start a discussion on Exec and auug-announce.

Action MP

Chapter Council Meeting - MP proposes to run this at the conference. Announcement from MW. 5pm on Thursday at the conference. EC get room for 20. Carried

MP queried what is happening with membership - EC

MEETING CLOSED: 5.10PM

NEXT MEETING:
TUESDAY 7 SEPTEMBER 1999
AT THE CONFERENCE

❖

AUUG'2000

THEME "OPEN SECURITY"

It's only 11 months away!

AUUG '2000 will be held on the Gold Coast from
24th to the 28th of July, 2000.

Put the date in your diary now and start thinking
of possible topics for papers.

The call for papers will be in the next AUUGN.

AUUG Local Chapter Meetings 1999

CITY	LOCATION	OTHER
BRISBANE	Inn on the Park 507 Coronation Drive Toowong	For further information, contact the QAUUG Executive Committee via email (qauug-exec@auug.org.au). The techno-logically deprived can contact Rick Stevenson on (07) 5578-8933. To subscribe to the QAUUG announcements mailing list, please send an e-mail message to: <majordomo@auug.org.au> containing the message "subscribe qauug <e-mail address>" in the e-mail body.
CANBERRA	Australian National University	
HOBART	University of Tasmania	
MELBOURNE	Various. For updated information See: http://www.vic.auug.org.au/auug/vic/av_meetings.html	The meetings alternate between Technical presentations in the odd numbered months and purely social occasions in the even numbered months. Some attempt is made to fit other AUUG activities into the schedule with minimum disruption.
PERTH	The Victoria League 276 Onslow Road Shenton Park	Meeting commences at 6.15pm
SYDNEY	The Wesley Centre Pitt Street Sydney 2000	

Up-to-date information is available by calling AUUG on 1800 625 655.

Unix Traps and Tricks

Jerry Vochtelo
jerry@socs.uts.edu.au

Welcome again to the UNIX tricks and traps column. There were no contributions received for this issue so for next issue, if you have anything please send it to me; jerry@socs.uts.edu.au, otherwise you have to be subjected to my dodgy scripting.



XTERMINAL CONTROL SEQUENCES

Jerry Vochtelo
jerry@socs.uts.edu.au

In the column this week I will talk a little about xterm control sequences. If you are not running X then I apologise in advance. The state and internal parameters of Xterminals can be changed by typing the appropriate keys in the terminal. For example:

```
echo -n "\[\[1mhello"
```

will print the word "hello" in a bold font.

I use the xterm control sequences for a couple of reasons.

This first script is called title and it changes the name of your xterm (as well as the icon name)

```
#!/bin/sh
#
# This will rename the xterm and the icon which this is run in
#
# Author J.Vochtelo 26/11/96
# (c) 1996, All rights reserved
#
# To do the the character controls use ^V^<key>
# or ^Q^<key> in emacs

if [ $# -eq 2 ]; then
    echo -n "\[\[2;$2^G^\[1;$1^G"
elif [ $# -eq 1 ]; then
    echo -n "\[\[0;$1^G"
else
    echo "Usage: title [icon] window"
fi
```

The next script changes your xterm colours (if your xterm supports colours):

```
#!/bin/sh
#
# This will change the colour of an xterm
# without any options set to white foreground and black bg
# with just one option, takes it to be the foreground colour
# otherwise fg and bg are set by the -f and -b flags respectively
#
#
# Author J.Vochtelo 1/2/99
# (c) 1999, All rights reserved
#
# To do the the controls use ^v^<key> (vi)
# ^q^<key> (emacs)
```

```

if [ $# -eq 0 ];
then
    echo -n "^[10;white^G";
    echo -n "^[11;black^G";
else

while ( [ $# -ne 0 ]; )
do
    if [ $# -gt 1 ];
    then
        if [ $1 = "-f" ];
        then
            shift;
            echo -n "^[10;$1^G";
            shift;
        else
            if [ $1 = "-b" ];
            then
                shift;
                echo -n "^[11;$1^G";
                shift;
            else
                echo "Usage: scol [fgcolour | [-f fgcolour] [-b bgcolour]]";
                echo "                -f colour      set foreground to colour";
                echo "                -b colour      set background to colour";
                break;
            fi
        fi
    else
        echo -n "^[10;$1^G";
        break;
    fi
done

fi

```

You can also do fun things like putting the control characters to change the title of an xterm into filenames. This will result in a name change when an `ls` is done on the directory (unless you do an `ls -b`, which we all do anyway, right? :). Similarly you can also make filenames appear in bold etc.

For those who are interested the full list of control sequences is at:

<http://www.mit.edu/afs/net.mit.edu/project/X11R5dev/src/hardcopy/clients/ctlseqs.PS.Z>

✧ ✧ ✧

HOW DO YOU FIND THE C: DRIVE IN LINUX?

David Newall
david.newall@tellurian.com.au

One of the questions which pesky DOS users will ask is how to find the C: drive in Linux. Perhaps the most concise answer is RTFM, which of course every true Unix Wizard binds as a macro to some easy to locate key such as the space-bar. Sometimes, if I'm in a particularly irksome mood, I like to answer thus...

The C drive is the first DOS or Windows partition. It's probably on your first IDE hard drive, so the Linux name will be `/dev/hda-something`. If the C drive is on the first partition it will be `/dev/hda1`; if it's the second partition it will be `/dev/hda2`; and so on. To "find" the C drive you must mount it. The following command will work:

```
# mount /dev/hda1 /
```

Following this your C drive will be available at `/`.

✧ ✧ ✧

UNIX IN MAINTENANCE MODE

David Newall
david.newall@tellurian.com.au

Ever had the problem that your most basic Unix utilities weren't available, such as `ls`, `cat` and `more`? Perhaps this was when you were rebuilding your system after a disaster. Install disk often permit you to use a shell, but due to space limits on a floppy, they only include the utilities that are absolutely essential.

You can build a surprising set of "standard utilities" using just the Bourne shell:

Listing directories

```
ls() {
  [ -z "$1" ] && echo * || echo $1/*
}
```

Displaying the content of files

```
cat_out() {
  while read cat_line; do
    echo "$cat_line"
  done
}

cat() {
  if [ -z "$1" ]; then
    cat_out
  else
    for cat_file; do
      [ "$cat_file" = - ] && cat_out || cat_out < $cat_file
    done
  fi
  unset cat_file cat_line
}
```


Paging through files

```
more_out() {
  more_count=0
  while read more_line; do
    [ $more_count -eq 23 ] && {
      echo "--more-- \c"; read more_count; more_count=0
    }
    more_count=`expr $more_count + 1`
    echo "$more_line"
  done
}

more() {
  if [ -z "$1" ]; then
    more_out
  else
    for more_file; do
      [ "$more_file" = - ] && more_out || more_out < $more_file
    done
  fi
  unset more_file more_line more_count
}
```

I might just say that with a little thought and effort a great many problems can be solved with this technique. For example, I recently was unable to log in to a machine. I guessed (correctly) that all process slots were full; and noticed that `getty` respawned when it quit. "Race," thought I, and proceeded to log in on one screen while quitting `getty` on another. With no free process slots it was a challenge to discover exactly what was wrong. For the standards conformant amongst us, probably `exec ps` is necessary, which means winning another race to log in again. Those who are not so standards challenged can `ls /proc` (see above). You can't use pipelines when there are no spare process slots, but you certainly can use temporary files to hold the output from each stage. A shell-coded, specialised `grep` resulted in an appropriate kill statement, and all that was left to do was `exec 0<kill-file` and presto! one working system.

P: AUUG Incorporated - The Australian UNIX and Open Systems Users Group - Netscape
File Edit View Go Command Help



AUUG Incorporated Web Home

Welcome to the AUUG Web Server. AUUG is the Australian UNIX and Open Systems User Group. AUUG is a national body, with chapters that organise local activities in most capital cities.

- [About AUUG](#)
- [Annual Conference](#)
- [Summer Conferences](#)
- [Tutorials and Road Shows](#)
- [John Lions Award](#) for work in Open Systems
- [AUUGN and Other Publications](#)
- [AUUG Local Chapters](#)

Home
Back Forward Stop

Check out the
latest
AUUG '99
Conference
News

And keep up to
date with the
happenings at
the AUUG
website:

www.auug.org.au

Application for Institutional Membership

Section A: MEMBER DETAILS

The primary contact holds the full member voting rights and two designated representatives will be given membership rates to AUUG activities including chapter activities. In addition to the primary and two representatives, additional representatives can be included at a rate of \$70 each. Please attach a separate sheet with details of all representatives to be included with your membership.

NAME OF ORGANISATION: _____

Primary Contact

Surname _____ First Name _____
 Title: _____ Position _____
 Address _____
 Suburb _____ State _____ Postcode _____
 Telephone: Business _____ Facsimile _____
 Email _____ Local Chapter Preference _____

Section B: MEMBERSHIP INFORMATION

Renewal/New Institutional Membership of AUUG \$350.00
 Surcharge for International Air Mail \$120.00
 Additional Representatives Number @ \$80.00

Rates valid as at 07/96

Section C: PAYMENT

Cheques to be made payable to **AUUG Inc** (Payment in Australian Dollars only)

For all overseas applications, a bank draft drawn on an Australian bank is required. Please do not send purchase orders.

-OR-

Please debit my credit card for A\$ _____

Bankcard Visa Mastercard

Name on Card _____
 Card Number _____
 Expiry Date _____
 Signature _____

Please mail completed form with payment to: _____ Or Fax to: _____

Reply Paid 66
 AUUG Membership Secretary
 PO Box 366
 KENSINGTON NSW 2033

AUUG Inc
 (02) 9332-4066

Section D: MAILING LISTS

AUUG mailing lists are sometimes made available to vendors. Please indicate whether you wish your name to be included on these lists:

Yes No

Section E: AGREEMENT

I/We agree that this membership will be subject to rules and by-laws of AUUG as in force from time to time, and that this membership will run from time of joining/renewal until the end of the calendar or financial year.

I/We understand that I/we will receive two copies of the AUUG newsletter, and may send two representatives to AUUG sponsored events at member rates, though I/we will have only one vote in AUUG elections, and other ballots as required.

Signed: _____

Title: _____

Date: _____

AUUG Secretariat Use

Chq: bank _____ bsb _____

A/C: _____ # _____

Date: _____ \$ _____

Initial: _____ Date Processed: _____

Membership#: _____



UNIX® AND OPEN SYSTEMS USERS

Membership
Application

AUUG Inc Secretariat

PO Box 366, Kensington NSW 2033, Australia

Tel: (02) 9361 5994

Free Call: 1 800 625 655

Fax: (02) 9332 4066

email: auug@auug.org.au

ACN A00 166 36N (incorporated in Victoria)

http://www.auug.org.au

AUUG Inc is the Australian UNIX and Open Systems User Group, providing users with relevant and practical information, services and education through co-operation among users.

AUUG OFFERS SOMETHING FOR YOU!!

AUUGN
Technical Newsletter
 AUUG's bi-monthly publication, keeping you up to date with the world of UNIX and open systems.

Education
 Tutorials
 Workshops

Events.....Events.....Events

- Annual Conference & Exhibition
- Overseas Speakers • Local Conferences
- Roadshows • Monthly meetings

DISCOUNTS
 to all AUUG events and education.
 Reciprocal arrangements with overseas affiliates.
 Discounts with various internet service providers, software, publications and more...!!

Connections

- Newsgroup aus.org.auug

Application for Individual or Student Membership

Section A: PERSONAL DETAILS

Surname _____ First Name _____
 Title: _____ Position _____
 Organisation _____
 Address _____
 Suburb _____ State _____ Postcode _____
 Telephone: Business _____ Private _____
 Facsimile: _____ E-mail _____

Section B: MEMBERSHIP INFORMATION

Please indicate whether you require Student or Individual Membership by ticking the appropriate box.

RENEWAL/NEW INDIVIDUAL MEMBERSHIP

Renewal/New Membership of AUUG \$100.00

RENEWAL/NEW STUDENT MEMBERSHIP

Renewal/New Membership of AUUG \$25.00
 (Please complete Section C)

SURCHARGE FOR INTERNATIONAL AIR MAIL \$60.00

Rates valid as at 07/96

Section C: STUDENT MEMBER CERTIFICATION

For those applying for Student Membership, this section is required to be completed by a member of the academic staff.

I hereby certify that the applicant on this form is a full time student and that the following details are correct.

NAME OF STUDENT: _____

INSTITUTION: _____

STUDENT NUMBER: _____

SIGNED: _____

NAME: _____

TITLE: _____

DATE: _____

Section D: LOCAL CHAPTER PREFERENCE

By default your closest local chapter will receive a percentage of your membership fee in support of local activities. Should you choose to elect another chapter to be the recipient please specify here:

Section E: MAILING LISTS

AUUG mailing lists are sometimes made available to vendors. Please indicate whether you wish your name to be included on these lists:

Yes No

Section F: PAYMENT

Cheques to be made payable to **AUUG Inc**
 (Payment in Australian Dollars only)

For all overseas applications, a bank draft drawn on an Australian bank is required. Please do not send purchase orders.

-OR-

Please debit my credit card for A\$ _____
 Bankcard Visa Mastercard

Name on Card _____
 Card Number _____
 Expiry Date _____
 Signature _____

Please mail completed form with payment to: _____ Or Fax to: _____
 Reply Paid 66 AUUG Inc
 AUUG Membership Secretary (02) 9332-4066
 PO Box 366
 KENSINGTON NSW 2033
 AUSTRALIA

Section G: AGREEMENT

I agree that this membership will be subject to rules and by-laws of AUUG as in force from time to time, and that this membership will run from time of joining/renewal until the end of the calendar or financial year.

Signed: _____
 Date: _____

AUUG Secretariat Use

Chq: bank _____ bsb _____
 A/C: _____ # _____
 Date: _____ \$ _____
 Initial: _____ Date Processed: _____
 Membership#: _____