

SAGE

Advice

The Journal of The System Administrators Guild of Australia

ISSN: 1447-5049

Volume 9 Number 1

March 2003 Edition

This Issue

SSH Tricks

Back up, say what?

Review of the 2003 Tasmanian IT Summer Conference

DNS Tracer

Data Center Management : Patching

Golden Sheep and Tiger Tails

SAGE-AU

The System Administrators Guild of Australia

<http://www.sage-au.org.au>

Contents

Executive

President's Report 3

Regional Groups

Regional Group contact information 20

General Information

Membership Drive 7
 Bulk Memberships 16
 Sponsorship 19

Articles

SSH Tricks 5
Brad Marshall
 Back up, say what? 6
Chris Herrmann
 2003 Tasmanian IT Summer Conference 8
Geoffrey Day
 DNS Tracer 11
Edwin Groothuis
 Data Center Management : Patching 14
Craig Armour
 Golden Sheep and Tiger Tails 18
Andrew Hennell

Advertising Rates (all prices are G.S.T. inclusive)

Inclusions with SAGE Advice

Slim Booklet or CD inclusion \$660
 (if more than slim booklet, \$poa)
 Single page inclusion \$385

Normal Rates

Full page, 170 x 250 mm \$275
 1/2 page, 170 x 120 mm \$165
 1/4 page, 170 x 55 mm \$88
 1/8 page, 80 x 55 mm \$44

Classified space

\$10 per ad, no more than 50 words

Bulk Rates

Ad space in four or more issues paid in advance attract a 20% discount.

See www.sage-au.org.au/newsletters_advert.html for our publication schedule for 2003.

Information

SAGE Advice is the official newsletter of the System Administrator's Guild of Australia. It is produced four times per year and distributed free of charge to members of SAGE-AU via snail mail.

Editor

Andrew Whyte
editor@sage-au.org.au

Membership & General Information

Lee Monette
lee@sage-au.org.au

Contact Details

Contact details for SAGE-AU are:

SAGE-AU
 346/500 Centre Road
 Bentleigh VIC 3204
 Phone: (03) 9822 7271
 Fax: (03) 9822 0260

Internet Server

<http://www.sage-au.org.au/>

SAGE-AU Committee

President

Andrew Hennell
andrew.hennell@member.sage-au.org.au

Vice President

Andrew Whyte
andrew.whyte@member.sage-au.org.au

Secretary

Simon Woodhead
simon.woodhead@member.sage-au.org.au

Treasurer

Steven Plunkett
steven.plunkett@member.sage-au.org.au

General Committee

Simon Coggins
simon.coggins@member.sage-au.org.au
 Rohan Durrant
rohan.durrant@member.sage-au.org.au
 Phil Kernick
phil.kernick@member.sage-au.org.au
 Brad Marshall
brad.marshall@member.sage-au.org.au

Associate Representative

Tim Oaks
timothy.oaks@member.sage-au.org.au

President's Report

Andrew Hennell

(president@sage-au.org.au)

Greetings, and welcome to another SAGE Advice newsletter.

Most of you by now have seen the new member website (<https://member.sage-au.org.au/>), and we've had plenty of positive feedback about it. I'd like to thank Andrew Whyte for his efforts over the past 6 months to prepare the site, program the back end, and get it all working.

The improved back end improves the way Lee manages our membership, and the front end allows you to view & edit your personal information with SAGE-AU, and manage your mailing list subscriptions.

But the job isn't over - the member site now allows the Guild to develop a range of member-only services, accessible through a single login from the one site. One such service is the 'knowledge project', to be developed over the coming months.

The knowledge project will provide a single access point to the SAGE-AU knowledge base, and will include:

- Articles from the newsletter
- Papers from our Annual Conferences
- Tutorial notes from our Annual Conferences
- Research papers
- Mailing list summaries
- Mailing list archives & searches

There is a lot of work to be done to bring this project to fruition - it's perhaps our largest project to date. Over the coming month I'll be seeking input and assistance from members keen to help with sorting and publishing the information available.

Our hope is that once the knowledge project is complete, SAGE-AU will have the most comprehensive single source of System Administration knowledge available within Australia. With your help we can achieve this goal.

Which leads me on to the Annual Conference - this year will be our largest yet, with 5 concurrent tutorial streams over three days, and two days of a content-packed technical conference.

For the first time in many years, the cost of the conference has risen to accommodate rises in every expense associated with running the conference. The increase was applied to the full price of tutorials and confer-

ence attendance, which means that when you factor in your SAGE-AU member and early bird discounts, the real rise is only a few percent.

This year we have content available for all levels of system administrator, and a subject range that should suit everyone. The draft tutorial program should be on the conference website by the time you read this newsletter, and the technical conference program will follow shortly.

The level of interest in the conference is already exceeding our expectations, so I urge all members to begin speaking with their employers to ensure they're available to attend, and to book as soon as registrations are announced.

The conference is being held in Hobart between the 4th & 8th of August - more details can be found at www.sage-au.org.au/conf/

Finally, to matters of chapters and regional meetings: For years the regional meeting has been the staple of SAGE-AU, allowing members to meet, to learn, and to share experiences that help them perform their jobs.

In Queensland, members around Rockhampton have begun regular meetings with great success, leading to an increase in membership in that area. Several other regions are also enjoying regular and successful meetings around the country.

A successful region is usually a reflection upon the state chapter exec, those who put their names forward to coordinate the chapter's meetings and increase membership within the state. I'd like to thank all chapter exec for their time and commitment to running local meetings - without their dedication these events simply wouldn't exist.

I'd also like to encourage members who'd like to assist in running meetings, be they in capital cities or any area where SAGE-AU members are, to put their hand up. Just as Rockhampton is now running successfully, many other regional centres may well benefit from local SAGE-AU meetings. If you're interested, talk with your chapter exec, or email me for more details.

Until next time, take care.

Andrew.



AusCERT Asia Pacific Information Technology Security Conference 2003

11th - 15th May
Royal Pines Resort
Gold Coast
Queensland

An international conference on IT security for CFOs, CIOs, CTOs and technical staff from industry, government and universities.

Visit: <http://conference.auscert.org.au/conf2003/>

Nominated by the delegates as "Australia's premier IT Security event"

Learn from world class experts about the latest strategies to make your information systems more secure and how to address computer security breaches. World class speakers will be present from Asia, Australia, Europe and the USA.



International speakers include: Gene Kim, Tripwire, Dave Dittrich, University of Washington, Wietse Venema, IBM T J Watson Research Center, Steve Hansen, Google, Inc, Jeff Carpenter, CERT/CC, Marcus Sachs, The White House, Simple Nomad, NMRC and BindView, Martin Khoo, SingCERT, Nick Fitzgerald, Computer Virus Consulting Ltd, Alex Shipp, MessageLabs, David Lynas, COSAC, Thomas Rude, Independent Consultant.

Local speakers include: John Barlow, AARNet, Andrew Clark, ISRC/QUT, Hamish Fraser, Oplus, Ajoy Ghosh, Independent Consultant, Karl Hanmore, Bank of Queensland, Gretchen Golik, Qantas, Adrian McCullagh, Freehills, Mark McPherson, AusCERT, Allison Stanfield, <e.law> Australia, Andrew van der Stock, b-sec, John Worthington, Business Continuity Institute of Australia, Matt Whelan, Independent Consultant, Steven Stroud, Scott MacLeod, Michael Cohen, David Collett, DSD.

To register: Early bird registration ends on 28th March. Special rates are also available to AusCERT members, CSIRTs and for multiple registrations.

Visit: http://conference.auscert.org.au/conf2003/registration_conference.html

Do you have all the information you need to protect your computer networks?

AusCERT is the national Computer Emergency Response Team for Australia and a leading CERT in the Asia/Pacific region. AusCERT maintains a world recognised reputation and trusted contact network of computer security experts around the world and provides prevention, response and mitigation strategies for members.

For information about AusCERT's membership packages and services call 3365 4417 or visit: <http://www.auscert.org.au/membership>

Training

SAGE-AU members receive 10% discount on non-AusCERT member prices. Currently offering:

- Introduction to Network IDS
- Network Monitoring for System Administrators

for further details visit: <http://www.auscert.org.au/training>

Independent Advice • Confidential Assistance • World Class Expertise



**Australia's National
 Computer Emergency
 Response Team**

SSH Tricks

Brad Marshall

(brad.marshall@member.sage-au.org.au)

There are some interesting things you can do with ssh, other than just getting a remote shell. These include a ssh gateway to access hosts behind a firewall, accessing cvs from behind a firewall, and administering embedded router systems securely.

SSH gateway

The SSH gateway works by forcing a ssh to another host as a particular username, based on the ssh key. The assumption is that there is a gateway box between the internal hosts and the internet, that has ssh running on it. Create a user on the gateway box, in this example user gw. Then edit the users authorized_keys file and add a line similar to the following:

```
command="ssh -t username@internal.host" \  
ssh-dss AAAAB3N . . . . 95nxu8Zjg username@host
```

To use the gateway, you then use something like:

```
$ ssh gw@firewall.example.com
```

This should, assuming you have the correct ssh key in the authorized keys file and the correct permissions, ask you to log into internal.host as the specified username.

CVS gateway

The CVS gateway is similar to the ssh gateway in that it forces an action based on the ssh key. However, there is also the requirement for ssh agent to be running and for agent forwarding to be enabled.

Create a user for this gateway, called cvs, then add the following to the authorized_keys file.

```
command="ssh -t bmarshal@hopper /usr/bin/cvs server" \  
ssh-dss ANzaalk94t12 . . lkanbLKD1jq8Zjg username@host
```

First start by running the following commands:

```
$ ssh-agent bash  
$ ssh-add
```

This will start ssh-agent and add in the users keys. Then run the following:

```
$ export CVS_RSH=ssh  
$ cvs -d :ext:cvs@firewall.example.com:/path/to/cvs/repo co module
```

This will allow you to check out the module stored in the repository at /path/to/cvs/repo.

Embedded Router Administration

There are many embedded routers available now that provide both telnet and web based administration. As telnet and http are not encrypted protocols and the routers do not usually have enough grunt to run encryption easily the administration interfaces are not suitable for being made available over the internet, yet it is useful to do so.

There is a way of administrating these boxes securely, but there is a bit of setup required. First, if the router is providing network translation for the LAN behind it, set the router to port forward 22 to an internal host, and make sure you can connect to it.

In this example we will assume that the router is 192.168.1.1, and the internal box is 192.168.1.2. To setup a port forward to allow administration via http, run the following:

```
$ ssh -L 8080:192.168.1.1:80 username@router
```

You can then connect to http://localhost:8080/ to talk to the remote router's web interface securely. To do the same for telnet, replace port 80 with 23.

Conclusion

These simple tricks using ssh have proved themselves very useful, and are in use almost every day. As you have seen, there is nothing complicated about their setup, and they can provide a very useful service to you. ●

Back up, say what?

Chris Herrmann, Far Edge Technology (chris@faredge.com.au)

Recently I had the pleasure of obtaining and installing a new SDLT 220 drive - a brand spanking new 110G native external backup drive. We acquired the drive for a client who were trying to fit 70G a night onto a DDS3 tape, which given Amanda's inability to span partitions across tapes was causing me heart palpitations. The juggling acts required to shift data around various partitions to make it all fit were not insignificant. If you're dealing with big systems, then this is probably all passé, however for a small business, generating up to 10G of new data per day was quite a big deal. They generate such large numbers as they're a graphic design house, and frequently have very large files moving around / being rendered / postscripted and so forth. This meant that a level 1 or 2 backup¹ was trying to do anything up to 10G a night. A DDS3 tape, which after formatting will only hold about 9.5G native just wasn't up to it...

So, they took the plunge and arranged to get the SDLT drive. We looked at a couple of models, namely the new Ultrium drives coming out which are faster than this drive (up to 30M / second if you get the 230 model), and smaller DLT drives.

We ended up choosing this one mainly because Quantum came through with a very good price on the drive. The other requirements were that it needed to fit the biggest partitions we're backing up onto one tape - currently 70G. We're not likely to cross the 70G partition size for a while as that's the biggest scsi drive in the system. The other drives are smaller (36G). The important thing wasn't necessarily that it could do a full backup of everything every night, but that it could do a full cycle over three days. If you're an Amanda user you'll be familiar with this, if you're not I'll explain a bit further on. The drive itself is a Quantum / Tandberg drive - I'm not sure who rebadged who :o)

We're running it externally, on a SCSI3 LVD connection, and the thing makes the old DDS3 look like a 386 - the average data transfer rate over 120G is about 15M / second sustained. We're now in the interesting position where the 4 servers it backs up are too slow for the drive. Combined, the servers spend about 10hour compressing data for the backups, which take about an hour to write to tape.

We used to have two servers backing up 4 systems onto two separate DDS3 drives (that is, there were 2 backup servers, 1 backup drive each...). With the larger tape, we have the luxury of using Amanda for one of it's strengths - multi-system backup. A cron job kicks off the server process, which in turn asks each of the Amanda client hosts to perform a dump. The server keeps track of which disks on which hosts last performed full dumps, and schedules dumps so that they fit within your *dumppcycle*². Each of the clients compresses the data, and streams it across to the backup server. The backup server has a 40G IDE drive which is used as a holding area for the backups, so that each of the four systems can simultaneously spool their compressed images. Once an image of a partition is finished it's written to tape, which is very efficient for the tape drive as it means that the tape drive is generally running continuously. Doing this means that the average tape-write speed stays high, and generates less wear and tear on the tapes. Using four systems to compress data is also obviously much faster than using one system to compress each one in turn. The systems range in speed from dual 333's to dual 1Ghz machines, and it's no surprise that the faster the better. A single processor 733 will compress 10G of data much faster (read several hours) than a dual 333. Interestingly, the tape writing process itself isn't very intensive - it uses between 1 - 2.5% processor time.

The current setup is nice in that it's very easy for us to expand it to include more systems, without having a high impact on the backup server. We could comfortably quadruple the number of systems and / or amount of data being backed up and it wouldn't take much longer each night for the backup job to run.

From a heterogeneous network point of view, the support of Windows isn't fantastic - you can back up Windows (smb) shares and it will do so reliably and quickly, however you can't generate a system snapshot. For us this isn't a problem as our Windows based servers are all application servers that hold no data. If a machine falls over we grab the ghost image, rebuild the machine and deploy the application from a CVS export - fast and easy to do. This wouldn't work for you if you needed to back up Exchange datastores or SQL Server databases.

For the moment, however, I'm a happy SysAdmin, my users are happy, and backups are running like a dream.

Chris Herrmann works by day maintaining systems, designing software and extracting order from chaos. By night, he works by a good novel, good food and a nice red wine.

¹ A level one backup is what's commonly referred to a Differential, Level 2 is incremental.

² A dump cycle is the cycle within which every disk, on every host has a full backup. Often this matches your *tapecycle*, which is the number of tapes you're rotating. In our case, the dump cycle is smaller than the tape cycle so that we have 2 full backups of each disk at any point in time. If we had more tapes, then this would stretch it out to 3 full backups and so forth, however this requires many more tapes. We don't do a full backup of every disk daily even though they'll fit on the tape, as this takes a long time and we want the backups to finish before people start work in the morning. ●



Broken or lost keyring?

Contact Lee in the Administration office (lee@sage-au.org.au) and she will arrange a replacement to be mailed to you and for your password to be re-set to the new keyring number.

A reminder that your default password is your keyring number so keep it close by at all times!

membership drive

The SAGE-AU definition of a System Administrator is "Anyone who maintains a computer not solely for their own use". At present we represent only a small proportion of the System Administration community.

More members not only means we have greater exposure within the industry, but with government bodies such as NOIE as well. This translates to greater acceptance of SAGE-AU, further membership discounts and benefits, and industry recognition of your membership to a professional guild. It also ensures we can keep our membership fees at a low, affordable level (about the lowest for any professional guild, and certainly within the IT industry).

For extra motivation to find new members, we've launched an incentive scheme that will run until the 30th of June, 2003. Basically, for every new member you refer to SAGE-AU, you go into the draw for a \$220.00 gift voucher. One winner is drawn from every twenty entries, giving you at least a 1 in 20 chance of winning!

There is no limit on entries - we want you to refer as many new members as possible - so if you refer five people you go into the draw five times.

Just make sure that your referred member puts your name in the box marked "Where did you hear about SAGE-AU?".

Winners will be announced as they are drawn, and the gift voucher may be redeemed for any SAGE-AU merchandise or with any of our SAGE-AU sponsors or supporters.

Think about it - if each member finds just one new member this year, SAGE-AU will double in size.

Jamie Gillespie from our Queensland Chapter won the most recent \$220 prize and elected to donate his money to the Brisbane branch of St John Ambulance (apparently warm SAGE-AU jackets are not in great demand in sunny Brisbane!). We (and St John Ambulance) thank Jamie for this kind gesture.

We're well into the next twenty referrals - will you be the next winner?

2003 Tasmanian IT Summer Conference

Geoffrey Day

(geoffrey.day@member.sage-au.org.au)

SAGE-AU thanks the 2003 Tasmanian IT Summer Conference Sponsors:



i n v e n t

<http://www.hp.com>



<http://www.oracle.com>



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



NEOTECK
Pty Ltd

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

The 14th annual Tasmanian IT Summer Conference was held on the Monday 24th February, 2003. For the first time this event was run solely by SAGE-AU.

Organising this year's conference seemed very easy in comparison to past events. All speakers were confirmed before Christmas, meaning I had a nice break without worrying about filling in spots. We were even so well organised we had an "emergency fill in" speaker set to present if we had any last minute issues (sincere thanks to John Dalton). Speakers were again a good mix of half local and half interstate people.

Organising the venue, registrations, food, accommodation, flights the list goes on and on and for the first time Lee Monette did all of this work for us, leaving us to concentrate on speakers. For the first time we could accept online credit card payments, making registrations much easier for many people. Having Lee do the "administration" for the conference made it a much easier and pleasant experience for the organisers; she even flew down to handle the administration on the day which meant that we actually got to see all the talks for the first time ever!

The conference was announced and a few people registered, then nothing. Deciding that bribery works well we re-sent the announcement, adding that the first 10 registrations would receive a free Tas Chapter SAGE-AU/AUUG polo shirt. This resulted in a general increase in registrations, gradually moving on until we had about 30 people registered. HP agreed to provide an iPAQ as the lucky door prize and SAGE-AU made the mistake of announcing jackets were available, so I begged for one as a second door prize. The opportunity to win seems to be an incentive and with about 10 days to go we had 55 registrations. A bit of last minute

advertising and to my surprise we ended up with 69 registrations on the Friday before the conference, the biggest attendance for the last few years.

Sponsorship was amazing this year !!!!!!! many many thanks to the sponsors:

- HP - major door prize of an iPAQ* H3970 Pocket PC
- Oracle - paying for the drinks session following the conference
- NeoTech - silver sponsorship
- Comstra - silver sponsorship

The venue (The Old Woolstore) was very good; excellent video projector and tiered seating meant everyone could easily see the speaker and the screen. Air-conditioning could have been a bit better as the outside temperature went well over 30 degrees (yes it can get warm in Hobart), and inside we were cold in the morning and too warm in the afternoon.

Morning, afternoon tea and lunch provided an excellent opportunity to talk and meet other people from the Tasmanian IT community.

Most people stayed for the excellent drinks session (thanks again to Oracle) where the bar was opened to anyone with a conference name tag. The conference closed at 5:15 and by 5:20 the bar was packed. I stayed to clean up the room a bit before making it to the bar and attempting to catch up with the early starters.

Following the drinks session the organisers and speakers went for dinner at T42, a bar/restaurant on the docks. A good time was had by all; drinks and dinner by the water on a warm summer evening in Hobart. A pleasant end to a very successful day.

Summer IT Conference Talks

RDESKTOP - Matthew Chapman

Matthew is the original author of RDESKTOP, a free, unix implementation of a client for Microsoft's Windows Terminal Services. Matthew covered the why, what and how of RDESKTOP. One of the most interesting things discussed was the debugging of the RDP protocol. The protocol is always encrypted, that is unless you use Windows NT and set the server and client locales to be France (due to some very interesting laws on encryption in France - this is fixed in Windows 2000/XP). Not content with reading French Matthew went on to write a RDP proxy to sit between the client and server so he could continue to debug in English. RDESKTOP is now a solid working Unix/Linux client for Windows Terminal Services and continues to be improved by Matthew and the RDESKTOP team.

SNORT a case study - Vlasti Broucek and Paul Turner

Vlasti's main research is in forensic computation, the study of the laws, evidence, techniques and technology relating to computing. In relation to forensic computing the question was put "are intrusion detection systems any use?" As a case study SNORT was installed and then the results examined. In summary the results of the study were depressing. Not enough data, incorrect data, false positives ... the list went on and on. Unless a lot of time and effort/expertise is put into an IDS, the resultant output is basically useless for Forensic computing. In a discussion with Vlasti after the conference we agreed that IDS' were useful for security but not for forensic computing.

Solaris 9 new features - Richard Smith

Richard kindly filled in for Tim Cook who could not make it at the last minute. You can tell that Richard is not in sales, as he provided a practical overview of the new functionality provided in Solaris 9. The interesting stuff, war stories etc, came over a few beers later but as he said, it's off the record.

Large Scale Backup Systems - John Dyer

John has been involved in researching backup systems for the Australian Bureau of Statistics and presented a good

overview of current backup technologies and architectures. John did not recommend any particular product emphasising that each situation is unique and needs to be addressed individually. He did point out that finding pricing for the "enterprise" backup systems can be worse than learning how to use them, one reseller basically saying "go buy it from someone else, we can't work out the charging".

TPAN - Andrew Bergman, Nick Grundy, Ben Short

For first time presenters these three did a fabulous job, explaining the concepts, history and plans for the Tasmanian Public Airwave Network. This is envisioned as a free, fast, private network using wireless technology, spanning Tasmania's major population centres. I hope some people present will be able to provide them support and help in getting the TPAN up and running.

25+ years of computing - Tony Grainger

Tony spoke of PDP11's, paper tape, core memory, 8inch floppy drives and many more things that most of us consider history. He has worked on everything from a PDP8 to a Cray Supercomputer. Unfortunately Tony ran out of time, presenting only about a 1/3 of what he had in the notes. We are looking forward to Tony presenting his entire talk at the next monthly meeting.

Cat 6 Deployment - Tony Kyne

Tony is a man on a mission!! Cat 6 should be able to handle multi gigabit speeds. After being stung by a very poor installation he is passionate about how to get things done the right way. He mentioned many aspects of cabling Cat 6 that had never occurred to me including:

- * at gigabit speeds the signal travels as a field in the plastic sheath
- * don't use plastic cable ties, they pinch the cable too much
- * don't put too many cables in a bunch, the weight compresses the bottom cable too much


He went on to talk about testing and how to obtain and interpret the test results. One of the interesting points was that in general 3% to 4% of all terminations were not adequate and need to be redone. This means there is a 97% chance of getting a good connection if they are redone. Tony was very knowledgeable about his subject, passionate and spoke well, with many attendees saying it was an excellent talk. ●



James Patterson, Co-organiser and Session Chair

Email editor@sage-au.org.au with your article ideas, short stories, book or product reviews or comments and suggestions - next edition closes 9 June 2003.

Guess what Yahoo!® uses to manage their servers?



"The Cyclades-TS Series of Console Access Servers provides the highest port density and security at a very competitive price. By using Linux as the embedded OS, it offers the flexibility required to manage our dynamic environment. The Cyclades-TS is a key element to help us keep our servers up and running." - **Pete Kumler**, *Manager of Site Operations, Yahoo! Inc.*

NEW! Device Server
Cyclades-TS100 • 1 port



Cyclades-TS Series

Console Access Server

- 1/4/8/16/32/48 RS-232 ports on 1U of rack space
- IP Filtering, RADIUS, and Secure Shell (SSHv2)
- Linux, FreeBSD, Sun, HP, and IBM compatible
- No unintentional breaks (Sun)
- Off-line data buffering



Request for your FREE CAS booklet at www.cyclades.com



www.au.cyclades.com

T: +61 7 3279 4320

F: +61 7 3279 4393

australia@cyclades.com



DNSTRACER - Exploring the DNS Infrastructure

Edwin Groothius (edwin@mavetju.org)

Quick DNS intro

The Domain Name Server system is a globally replicated and distributed database which primary translate hostnames (www.sage-au.org.au) into IP addresses (66.216.68.159), route mail (@sage-au.org.au) to mail-hubs (sagemx.sage-au.org.au) and converts IP addresses (66.216.68.159) into hostnames (platypus.instaweb.com.au). Without it, we would have to use remember the IP addresses of the servers we want to connect to (telnet 131.155.132.36 4000) and it would be very hard to send emails as easy as it goes today (mcvax!moskvax!kremvax!chernenko).

Normally you don't have to worry about DNS, you just get the settings for the nameserver you have to use via PPP when dialling into an ISP or via DHCP when connecting to a LAN at a company. They make sure that their nameservers know where to get the rest of their data, which are initially the root-nameservers.

The root-nameservers are the 13 (13 logical, but physical more) most important nameservers on the Internet. They know where the rest of the DNS servers can be found.

Furthermore you have master and slave servers for a domain: the data for a domain is only manually changed at the master, the slaves transfer the data via the internal DNS mechanics.

Quick DNS example

If you're requesting the IP address of www.sage-au.org.au your nameserver will ask one of the root-servers for it. It will reply that it doesn't know it, but that the answer can be found at the DNS servers for .au and supplies a list with them and their IP addresses (The list is known as Authority Data, the IP addresses are known as Additional Data). Your server will ask the question again at one of the servers responsible for .au and get a similar answer: it doesn't know it, but it hands you a list of servers for .org.au and their IP addresses. This goes on until you're at the servers which are responsible for sage-au.org.au, in which case you get the IP address of www.sage-au.org.au (Answer Data).

Your server now caches the data for .au, .org.au, .sage-au.org.au and www.sage-au.org.au for a short time (the Time To Live) so that following requests for that data doesn't need to explore so much, it just can do a quick lookup of in it's own cache and returns the answer.

SPOF?

The DNS system is not really a SPOF, it is designed as a globally replicated and distributed database which means that if you can't reach one of the servers, you can try it at a different one. As there are 13 root-servers which know where to find the rest, there are 6 servers for the .au domain (6 logical with a total of at least 8 IP addresses), there are 9 servers for the .org.au domain and two servers for the sage-au.org.au domain. The location of the servers on the Internet and replication is used to overcome connectivity problems. Regarding the network, there isn't much which can go wrong. Regarding the administrative side of it, that's where things go wrong.

APOF!

When you register a new domain, you are asked what the nameservers are and if necessary also the IP addresses. Furthermore, these nameservers have to be configured to answer requests for that new domain and to exchange information between them. And actually data has to be served on that domain. Five places for things to go wrong!

Unknown nameservers

At the time of writing, one of the domains of a nameserver for .org.au has expired (for people interested: optus.net has expired at December 16th 2002 and after half a month it still hasn't been re-registered). That means that the IP address of the nameserver audns01.syd.optus.net can't be found and that this server will never be queried (after all, if you don't know an IP address you can't connect to it)

Wrong IP addresses

Changing the IP address of a nameserver is a pain and often it will be forgotten on one or two machines (Remember that switch in the cupboard which got installed a long time ago? Yes, that one too has the IP address of the DNS server hardcoded). Or that the registrar makes it impossible to change the IP address of the nameserver via their website because of all kind of internal checks.

Lame and stealth servers

Lame servers are servers which are mentioned in the NS records for a domain but are not authoritative for that domain. This can happen because of a typo in the IP address or a change which has never been fully finished (new server

added while it wasn't ready or old server data removed but never from the NS records).

Stealth servers are servers which are not mentioned in the NS records but are authoritative for that domains. For example servers which have been removed from the NS records but the configuration of the server never updated.

Old data on a server

When data is changed on the master server, the slaves will have to transfer it from there. But sometimes they can't because the master has disabled it for some reason. In that case the data on the slaves will get more and more obsolete.

Wrong data on a server

DNS server software has strange habits and one of them is often that if you end a name without a dot, it will add the current domainname to it. So if you see a zonefile with `www.sage-au.org.au.sage-au.org.au`, you know that they forgot to end it with a dot at the end.

Now what is dnstracer?

Remember the `traceroute(8)` utility? It shows the path an IP packet takes when you send it to its destination IP address. Remember `ntptrace(8)`? It shows the path of NTP servers which your NTP client is syncing on. `Dnstracer` is something similar, it shows you where a DNS server will go for its information. So if you want to know the path to `www.sage-au.org.au`:

```
[~] edwin@k7>dnstracer -s . -o www.sage-au.org.au
Tracing to www.sage-au.org.au via A.ROOT-SERVERS.NET, timeout 15 seconds
A.ROOT-SERVERS.NET [.] (198.41.0.4)
|\___ SEC3.APNIC.NET [au] (202.12.28.140)
|   |\___ ns3.melbourneit.com [org.au] (203.27.227.10)
|   |   |\___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) Got authoritative
answer
|   |       \___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) Got authoritative an-
swer
|   |\___ ns3.ausregistry.net [org.au] (203.18.56.43)
|   |   |\___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |       \___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |\___ ns2.ausregistry.net [org.au] (203.18.56.42)
|   |   |\___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |       \___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |\___ ns1.ausregistry.net [org.au] (203.18.56.41)
|   |   |\___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |       \___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |\___ audns01.syd.optus.net [org.au] (No IP address)
|   |\___ au2ld.csiro.au [org.au] (130.116.2.21)
|   |   |\___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |       \___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |\___ dns1.telstra.net [org.au] (203.50.5.200)
|   |   |\___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |       \___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |\___ box2.aunic.net [org.au] (203.202.150.20)
|   |   |\___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |       \___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|   |       \___ ns4.ausregistry.net [org.au] (210.8.15.253)
|   |           |\___ ns2.sage-au.org.au [sage-au.org.au] (130.102.171.100) (cached)
|   |               \___ ns1.sage-au.org.au [sage-au.org.au] (203.27.221.52) (cached)
|\___ SEC1.APNIC.NET [au] (202.12.29.59)
|   |\___ au2ld.csiro.au [org.au] (130.116.2.21) (cached)
|   |\___ dns1.telstra.net [org.au] (203.50.5.200) (cached)
|   |\___ box2.aunic.net [org.au] (203.202.150.20) (cached)
[...]
```

<code>ns1.sage-au.org.au (203.27.221.52)</code>	<code>www.sage-au.org.au -> 66.216.68.159</code>
<code>ns2.sage-au.org.au (130.102.171.100)</code>	<code>www.sage-au.org.au -> 66.216.68.159</code>

Just like expected: the server goes to a root-server, the servers for the .au domain, the servers for the .org.au domain and the servers of the .sage-au.org.au domains. The answers received are printed at the end and they all seem to agree on it.

Sometimes it will go wrong, for example when a server is unreachable or when a lame server is detected:

```
[~] edwin@k7>dnstracer -o -s RELAY-1.FTEL.CO.UK fataldimensions.nl.eu.org
Tracing to fataldimensions.nl.eu.org via RELAY-1.FTEL.CO.UK, timeout 15 seconds
RELAY-1.FTEL.CO.UK (192.65.220.24)
| \___ ns.cistron.nl [nl.eu.org] (62.216.31.55) Got answer
| \___ ns.lf.net [nl.eu.org] (212.9.160.1) Got answer
| \___ ns.eu.org [nl.eu.org] (137.194.2.218) Lame server
| \___ ns2.ispi.net [nl.eu.org] (206.131.193.15) Got authoritative answer
| \___ ns.patriots.net [nl.eu.org] (206.131.200.40) * * *
| \___ auth1.dns.elm.net [nl.eu.org] (81.17.34.251) Got authoritative answer
[...]
```

The difference between "Got answer" and "Got authoritative answer" is that the first one can be a cached answer, while the second one is one from a server which admits that it's responsible for that domain.

What can you see with dnstracer?

Dnstracer shows you the path from the root DNS servers to the DNS servers responsible for a domain. It shows if there are unreachable servers, lame servers but doesn't show servers which aren't configured for that domain. It will query for MX records, SOA records, NS records, normal A/AAAA and PTR records (and other ones). And at the end, it will print the results received. But it will not interpret the results for you.

More information

See <http://www.mavetju.org/unix/dnstracer.php> for more information about the dnstracer utility and how to obtain it. For FreeBSD and OpenBSD, it is in the ports-collection. For Linux, there is an RPM for it. Otherwise, just grab the tarball and compile it. ●

The System Administrator's Guild of Australia thanks our
generous Sponsor for their support.



Data Center Management: Patching!

Craig Armour

(craig.armour@ausbit.com.au)

Patching...

Leaving a system running with known bugs and security holes can be disastrous. By regularly and frequently patching a system, you will reduce the chance of any of these problems affecting you. Unfortunately patches aren't always perfect when they are released and can sometimes introduce bugs of their own. In this chapter we will discuss strategies you can implement in order to gain the best results out of your patch procedure.

The Basics

As a general concept, a complete patch procedure should include the following steps

- determine what patches need to be applied
- collate the necessary patches
- test all the patches on a suitable test environment
- obtain a full backup of your system
- apply the patches to your system
- reboot your system
- obtain a full backup of the newly patched system

No doubt, by now you are probably thinking "Do I really need to do all that to apply one patch?" But before we answer that, we should break each of the above points down into their respective elements.

Obtain Patch Lists

This is an obvious necessity as without a patch list, you are not going to know what patches the system requires. If you have access to any patch analysis tools such as patchdiag or patchpro, you can use these to determine the current state of your system and obtain patch lists in that way. Alternatively, there should be a list of recommended patches for your Operating System release. Either way, the number of patches and their revisions will change on a regular basis. It is important to regularly check for changes or additions.

Obtain Patches

Once you know what patches you require, you obviously need to get them. You might be able to get them from the Internet, or you might have a cd containing all the patches for your Operating System. You will need to find a method which best suits you.

Test the Patches

It is important that you develop a suitable patch test procedure. You should *NEVER* apply a patch to a mission critical system without first checking the patch will have no dire consequences for your system. Each patch should be released with documentation detailing what files the patch affects. Special install instructions may also be listed in this documentation. For this reason, these files should be reviewed as part of your testing procedure. We will discuss test strategies later.

Patch Application

How you intend to apply your patches may affect your test procedure. You can either apply each patch individually or as a cluster, but you should apply them to your production environment in the same way as you did to your test environment.

Reboot the System

This step is very often questioned. The whole purpose of applying patches is to improve availability, so wouldn't rebooting the system reduce the level of availability? Yes and no. Obviously rebooting the system is going to reduce your immediate level of uptime, however if the system fails to reboot now, it would be easier to associate any problems to the patching, then it would if the system was to fail in six months time.

Have a Full backup

Before making any changes, you should make sure you have a full backup of the system to tape. This will give you a means to revert your system to its prior form should the worst happen. Once it is verified that the system is running fine, another full backup of the system should be dumped to tape. By repeating this procedure each time a change to your system is made, you will always have a backup of the system to fall back on. Perhaps you should look at rolling this into your regular backup strategy.

So it can be seen that each of those steps are necessary for an effective patch procedure. By skipping any one of those steps, you greatly reduce the effectiveness of your procedure, and may even place your data and systems at risk. How you implement these steps will vary from enterprise to enterprise, however a general

strategy can be put into place and customized for your environment.

Strategies, The 3 Tier System

Generally, most sites can break their machines into three tiers, or categories. Each tier defines the importance of the system to the running of your organization. For the lack of a better naming scheme, we'll define them as Tier one, two and three.

- **Tier One:** These systems are the most mission critical systems to your organization
- **Tier Two:** Tier two systems are still important machines, but their availability is not vital to the running of your organization
- **Tier Three:** Tier three systems are generally test or development machines which are not in any way relied upon for their availability

By ranking your systems in this manner, it not only becomes easier to decide on an appropriate maintenance schedule for your systems, but it becomes much clearer in which order your machines should be patched.

Before we delve into how this system works, you need to ensure that the systems that exist at the tier one level, are replicated in some fashion at either the tier two, or the tier three level. That is, there are systems at the tier two and three levels that have the same combination of Operating System release and your Application release. If this can not be achieved, then the value of this system may be greatly reduced. The reasons for this will become clearer shortly.

The way this strategy works is simple. Patches are tested on release on the lower tier systems (tier three) and are gradually rolled up the tiers as the systems are brought down for scheduled maintenance. Where this system can start to become complicated is deciding on how often to patch the higher tier systems, and what patches to apply to these systems.

For the purpose of this example, we will assume that tier one systems are brought down once every six months for scheduled maintenance and tier two systems every three months. By the very definition of tier three systems, they can be rebooted at any time so scheduling maintenance windows is not so much of an issue. You will need to decide on an appropriate maintenance schedule for your enterprise. You may even wish to break your systems down into a higher number of tier levels. There is no fixed answer to what you should do here, you will have to decide what best suits you and even change things as time rolls on.

Which patches you apply will come out of your patch testing procedure. As a patch is released from your

vendor, apply and test the patch on your tier three systems using the steps mentioned earlier. Any immediate problems stemming from that patch should appear here. If the patch seems fine and your systems run fine then add that patch to your patch cluster for the month. At the end of the month or when the first maintenance window rolls around for your tier two systems, you can apply the cluster you've created for that month (and every other month's cluster that's not been already applied to that particular machine) to your system. Again, you must use the patch application steps mentioned previously paying particular attention to running a full backup of your system. This procedure is repeated until the first maintenance window for your tier one systems rolls around. By running your patch clusters on your tier two systems over a period of time , any problems that arise to your application, or to the system in general as a result of the patches should become apparent. If it is deemed that a particular patch has caused a problem, the patch should be removed from your systems, the relevant monthly patch cluster and your vendor should be contacted. If there are no problems it should now be considered safe to roll those patch clusters onto your tier one systems. Again, following proper patch procedure, apply all the previous months' patch clusters onto your system taking note that only patch clusters tested on lower tier systems are applied and again, make sure a full backup of the system is run before and after the system is patched. By rolling the patches up the levels of importance in this manner, any chances of a patch adversely affecting the availability of your mission critical systems is greatly reduced. Using this method should also help manage the patch levels of your systems. By simply recording the latest patch cluster applied to that system into a table or spreadsheet, it can be easily determined what patches need to be applied to a particular system during it's next maintenance window.

Caveats

There are a few things that you should think about before applying any strategy, however here are a few things to think about with this system.

System Redundancy

In order for this strategy to work, the fairly large assumption that all your systems are allocated regular scheduled down time, is made. In order for this to happen, There must be some level of system redundancy so that services are not completely interrupted while mission critical systems are down. A method for running your applications on alternate servers (such as clustering) should be devised. This not only provides

you with an avenue to shut down primary systems for regular, scheduled maintenance, but also provides a level of redundancy in the event a primary system fails.

Configuration replication

This strategy also makes the assumption that you have systems in the lower tier levels which replicate the configuration of systems in the higher tier levels. That is, you have systems in tiers two and three that have the same release of the operating system and the application as your systems in tier one. If this is not the case, you should try to replicate the system at the development level, even if it is only replicated for the purpose of testing. You may be able to solve this problem if you run your primary systems in a cluster environment. In which case you should only patch up one node at a time, making sure appropriate patch test procedures are used, and only patching up the other nodes when it is apparent that no problems have occurred due to the patches.

Patch consistency and Emergency Patches

All your systems should be at a consistent patch level before this strategy is implemented. If not, then impor-

tant patches could be missed, and managing each systems patch level can become more complicated. In the event that a patch needs to be applied urgently, a test procedure still needs to be put into place containing the steps mentioned earlier in section 2.1 . This procedure may vary depending on the level of urgency and the importance of the machine to your enterprise. However, the system should be monitored after such a patch has been applied, and record that the system varies from the rest in your records

Conclusion

So as you can see, the effort required in keeping the patch levels of your systems up to date is substantial. That being said, the consequences to reliability and availability of not keeping your patch levels up can be even greater. By developing an appropriate strategy such as the one detailed earlier, you can increase the effectiveness of your patch procedures, while reducing the over all complexity. ●

Do you have five or more SAGE-AU members at your workplace?

Contact Lee (lee@sage-au.org.au) to discuss how we can use the new bulk membership scheme to make your life (and that of your accounts department) that little bit easier with bulk renewals.

We also offer a small discount on entry fees for ten new members joining at the same time from the same workplace - contact Lee for further details. And remember, if they all put your name down as a reference, you go into the draw ten times for the \$220 prize (see pg 7).

Reproduced with kind permission of user Friendly

USER FRIENDLY by ILLIAD



www.userfriendly.org

Confernece website: www.sage-au.org.au/conf/

SAGE-AU

11th Annual Conference, Hobart 2003

Pete Calvert - Migrating NT4 to Windows 2000
 Brad Marshall - LDAP
 Elizabeth Zwicky - Backup systems
 Lee Damon / Mike Ciavarella - Effective System Administration
 Randal Schwartz - Learning more perl
 Pete Calvert - securing Windows 2000
 Phil Saunders - Storage 101
 Paul Yarrow - Active Directory 101
 Lee Damon - Issues in Unix infrastructure design
 Phil Saunders - SAN design and troubleshooting
 Elizabeth Zwicky - Think like a sysadmin
 Karl Hanmore - Outsourcing
 Andrew Radke - security practices for external networks
 Adrian Cockcroft (Sun) - Capacity planning
 Lee Damon - Ethics and professional conduct
 Dan Yarrow - Active Directory 101
 Adrian Cockcroft - Performance tuning
 Steven Penner - Tivoli
 Jennifer Walbank (Apple) - Apple osX
 Geoff Halprin - Exchange 101
 Dan Yarrow - Active Directory
 Phil Kernick - IPsec VPN's on Cisco IOS
 Jason Andrade - Storage
 Adrian Cockcroft - Solaris
 Pete Calvert - Certification & Degrees
 Gregory Bond - MRTG
 Baden Hughes - Network Monitoring
 Karl Hanmore - outsourcing
 Paul Lawrie - Cisco IOS networking
 Brian Meilak - Securing Cisco
 Zoong Pham - clustering Alpha's
 David Burren - Photography
 Tongue U (Sun) - Open source from Sun's perspective
 Gregory Bond - Email virus scanning

Where will you be in August 2003?

Join System Administration's best at the SAGE-AU Annual Conference

- 3 days of tutorials
- 2 days of technical conference
- Annual guild dinner
- Trivia night
- Birds of a Feather sessions
- Pre and Post conference tours
- Partner's travel program
- Discount accommodation

Contact:

Lee Monette
 SAGE-AU Operations Manager
 Email: lee@sage-au.org.au
 Phone: 03 9822 7271
 Fax: 03 9822 0260
 SAGE-AU
 346/500 Centre Road
 Bentleigh, VIC 3204

Golden Sheep & Tiger Tails

Andrew Hennell

(andrew.hennell@member.sage-au.org.au)

So I'm sitting at my keyboard trying to flesh out an article for the SAGE Advice newsletter. Why? I'm not 100% sure. Perhaps it's a vain streak in me coming to the surface, wanting to see my words published. Perhaps it's to keep the newsletter editor off my back. Or perhaps it's just to avoid attending to the growing jungle outside that is my garden. I'm tempted to believe it's more of the latter.

But what to write about? SAGE-AU members are technical; and system administration is a very technical profession, so I guess I should write about something technical. In a world surrounded by 'thingamajigs' and 'computer stuff', that should be easy enough... Or so you'd think.

Recently I've been pulled over to what many consider 'the other side' – the world that comprises management, consultants, contractors, and clients, considered by some to be worst of all. From the other side, I saw system administration as a service industry.

Not the service industry that greets you with a smile and asks if you'd like fries with that, nor the service industry that makes you go through IVR Hell, before placing you on hold only to drop the call after several minutes.

No, I'm talking about real service. The kind you used to get at service stations, when golden sheep were around and it was fashionable to stick a tiger tail out of your petrol tank. For those who don't remember such icons, Golden Fleece was a huge petrol station chain, and Esso promoted their fuel with 'get a tiger in your tank'.

Sadly, I was not old enough to drive when these things were about, but I do remember sitting in the back seat of mum's car and watching real service in action.

"I'll just check your oil ma'am," the operator would say to my mum while the wheels clicked over on the petrol pump. "Battery water's a bit low, I'll top it up for you" he'd holler from the mystical depths below the hood, only emerging to clean the windscreen and holster the petrol nozzle back into the pump.

Nothing was a problem, and I never saw the mechanic without a smile. My knowledge of engines and what went on under the hood of a car at that time was limited; my mother's was even less. It was the mechanic's job to keep it humming, and not to bother us with the technical mumbo-jumbo of those otherwise mystical depths.

Mum didn't care what grade oil the car took, or if it took oil at all. She wanted a car that was full of fuel and ready to work.

Turn the key and off you go. No fuss, no technical stuff, just make it work. And it did. All I wanted was a Kool Fruit from the tin mum had bought whilst paying for the petrol.

By now, you're all wondering where this is going. Has Andrew finally flipped? Well that's debatable, but I assure you the trip down memory lane has a purpose, if only to make me crave a Kool Fruit.

I believe that in system administration, everyone who uses your systems is your customer. I'll say that again, just for effect – everyone who uses your systems is your customer. It's your job to serve them with a smile, to be friendly and approachable, to anticipate problems, to smooth things over and to work occasional miracles.

I can sense the apprehension already. Us? Be nice? Surely not! Long live the BOFH! Well my friends, the days of the BOFH are gone and, I suspect, it's about time.

Helping customers is a rewarding process. They come to you for help, because you're the recognised expert in your field. You can help them. That gives you a lot of importance in this equation, and you can use that to your advantage or detriment.

If contempt breeds contempt, then the converse is true – happiness breeds happiness. When you help a user solve their problem quickly, do you feel good or bad?

Did they smile and thank you for your help? Or did they hump and walk off because it had taken them half a day or more to get you to help?

Handling users' requests is what we do. Ensuring the systems are running 100% is what we're here for. System Administration is service.

Nobody snarled at the driveway mechanic working in the shadow of the large golden sheep. Nobody felt bad about asking him for help, and he never begrudged offering that help. Always available, always friendly, always there for the clueless driver.

He never frowned and thought "idiot driver, hasn't kept his oil levels right" – nor did he think "you're the driver, clean your own windscreen".

No, that mechanic knew what service was. He knew that the driver wanted to drive, and not much else. He knew it was his job to provide the service that gave the drivers what they wanted.

Occasionally I hear horror stories about dealing with compu-

ter support personnel and system administrators. Whenever I hear of these bad experiences, I cringe – not only for system administration as a profession, but for society's loss of service.

Next time you're asked to help a user, think positively. The user has expectations of the service you'll provide – often it's

as simple as getting them back online. Don't just reach their expectations – exceed them. Check their oil and water (ok, perhaps their disk space and patches), give service with a smile, and see how your day can improve.

Perhaps one day, we'll see golden sheep and tiger tails again. For now, I'm off to buy a pack of Kool Fruits. ●

The Rewards of

Gold Sponsorship

Gold Sponsors of SAGE-AU have the benefits of:

- Full page ad in every edition of SAGE Advice and listed in Sponsor Section
 - Gold listing at SAGE-AU Conferences
 - Listing on all publications
 - Listing on all SAGE-AU www pages with logo and link
- Listing as Gold on SAGE-AU's sponsorship and conference pages with logo and link
 - Discounted SAGE-AU event attendance
 - Banner space at the conference, logo displayed on sponsor slide and acknowledgement at the start of each session.
 - Any provided literature displayed on a sponsors and information table.
- When offered, 50% discount on exhibition space at the annual conference.
 - Provided literature distributed in conference showbag.

If you would like a space like this reserved for your business contact SAGE-AU on (03) 9822 7271 or at secretary@sage-au.org.au

We also have other sponsorship levels available (see www.sage-au.org.au/sponsor.html) or can tailor a sponsorship package to suit your particular requirements.

Your company may also be interested in participating in our supporters programme (see www.sage-au.org.au/memberservices.html) whereby our members are offered discounts on goods or services upon presentation of their SAGE-AU member card.

Contact lee@sage-au.org.au for further information.

Regional Groups

Full lists of Chapter executive committee members, meeting schedules and up to date Chapter information can be found at <http://www.sage-au.org.au/regional.html>

Victorian Chapter

The Victorian group currently meets on the second Tuesday of the month at 7:00pm at:

VTR Consulting
Roseneath Place, South Melbourne

President

Chris Williams
christopher.williams@member.sage-au.org.au

Secretary

Thomas van de Hoef
thomas.van.de.hoef@member.sage-au.org.au

Treasurer

Ben White
ben.white@member.sage-au.org.au

The group mailing list is sage-vic@sage-au.org.au

Queensland Chapter

The Queensland group currently meets on the second Thursday of the month at 7.00pm at:

Room 102
Frank White Building (No. 43)
Cooper Road
University of Queensland, St Lucia

President

Mark Suter
mark.suter@member.sage-au.org.au

Secretary

Jamie Gillespie
jamie.gillespie@member.sage-au.org.au

Treasurer

James Lever
james.lever@member.sage-au.org.au

The group mailing list is sage-qld@sage-au.org.au

There is also a **Rockhampton** group - see

www.sage-au.org.au/rgql.html

South Australian Chapter

The South Australian group currently meets on the last Tuesday of the month at 6.30pm at:

Vectra Corporation
38 Sir Donald Bradman Drive, Mile End

President

Phil Kernick
phil.kernick@member.sage-au.org.au

Secretary

Monica Ciric
monica.ciric@member.sage-au.org.au

Treasurer

Grant Ward
grant.ward@member.sage-au.org.au

The group mailing list is sage-sa@sage-au.org.au

New South Wales Chapter

The New South Wales group currently meets in Sydney and Parramatta.

Acting President

Neil Glennon
neil.glennon@member.sage-au.org.au

West Sydney Organiser

Francis Liu
francis.liu@member.sage-au.org.au

The group mailing list is sage-nsw@sage-au.org.au

Australian Capital Territory Chapter

The ACT group currently meets on the last Tuesday of the month (except December) at 6.00pm at:

University House
Cnr Balmain Cres & Liversidge St, Acton

Interim President

Roy Meuronen
roy.meuronen@member.sage-au.org.au

Interim Secretary/Treasurer

Steve Jenkin
steve.jenkin@member.sage-au.org.au

The group mailing list is sage-act@sage-au.org.au

Tasmanian Chapter

The Tasmanian group currently meets on the third Monday of the month at a different location depending on the event being held (check <http://www.sage-au.org.au/rg/tas> for info).

President

Chris Fishwick
chris.fishwick@member.sage-au.org.au

Secretary

Andreas Modinger
andreas.modinger@member.sage-au.org.au

Treasurer

Anthony Vialle
anthony.vialle@member.sage-au.org.au

The group mailing list is sage-tas@sage-au.org.au

West Australian Chapter

The West Australian group currently meets on the first Wednesday of the month at 6.30pm at:

Moon & Sixpence British Pub
300 Murray Street, Perth

The group mailing list is sage-wa@sage-au.org.au

Please support your local SAGE-AU Chapter : join the Chapter executive committee, start a new regional group, offer your office boardroom as a venue, give a talk on a subject of interest, or simply turn up and be involved. Contact the group mailing list or sage-au-exec@sage-au.org.au to see how you can help.