

Saturday, February 23, 2019

The 'A Frame' for Managers

I have this analogy that managing people is like standing on the back of a moving truck trying to hold on to several sheets of lightly oiled glass. No matter how good a grip you might have, it never seems enough, and one good bump and any one of those fragile things can slip and shatter.

Every occupation has its 'tools of the trade'. Carpentry can't take place without a hammer, nor plumbing without a stilson. Yet, it seems to me that in many cases someone is simply 'made' a manager by the bestowing of that title, and left to cope bare handed with the situation.

If there is such a person as a 'natural' manager, I have never met them. Just like a runner can't naturally win a gold medal, it requires a lot of commitment, planning and support to get to the level where a win is possible.

Posted by Steve Waddington at 11:45

Friday, April 12. 2013

More NBN Mis-Information

It makes me angry every time I see the NBN ads on TV; the implications the average 'uninformed' view is drawn to conclude about how great and fantastic the NBN will be. Good job by the advertising agency to fulfill its brief, disgraceful job by a government to give such a brief.

I was asked to comment on the coalitions NBN alternative the other day. Part of which appeared here.

Here is the full response:

Really, I despair. But ok, here is my view on the main points:

1. Deliver 25Mbps by the end of the first term.

Why is it stupid:

Right now, ADSL2 delivers up to 24Mbps. So the promise is a 1Mbps increase, for only 20 Billion dollars. But without spending 1 cent, the current broadband service providers would have, if the NBN had not killed off all investment, reached that mark in the not too distant future anyway.

Why it is better than the current plan:

Spending \$20 Billion is better than Labors plan to spend \$40 Billion+++ , for still, essentially what we already have.

2. Fibre to the Node vs Fibre to the House

The statement "unlike the government's fibre-to-the-premises technology, the coalition's network would not be easily upgradable." is wrong.

The fact is, copper is a cheaper technology to deploy for short distance signaling, and even has advantages over fibre for ease and flexibility of installation. It is also easier and cheaper to replace a copper based network switch than to replace or upgrade a fibre switch.

Ask yourself, is your desktop or office server connected by fibre? No. It is connected by copper wire. To say that high speed data mandates fibre, at least for short distances, simply ignores physics, as well as the evidence in front of your eyes.

So fibre to the house, much reducing the roll out cost, is a point in favor I would say, if it weren't for the pointlessness of it all re: points 1 and 3.

3. Completed in 2019

Since it is 96% completed now - 24Mbps on ADSL2 vs 25Mbps, it is hard to see how that promise couldn't be kept.

4. Sell the network once it is ready

Well, we, taxpayers, already paid for the first network - PMG/Telecom/Telstra, which has been partly sold. Now we are paying for a second network, which will be sold. So I guess we can all look forward to a reimbursement check from the ATO once the sale goes through? Excuse my cynicism.

5. "opposition claims the government's NBN plan could blowout to \$90 billion - a claim the government venomously denies."

And well they may deny it. \$150bn is a more likely figure.

6. 25/50/100Mbps

In this blog post in 2009 I calculated the cost of a 1Gbps home user broadband service to be \$8bn, so even allowing for inflation, \$20bn still seems to high by a factor of 2:

<http://steve.blogs.exetel.com.au/index.php?/archives/227-1Gbps-to-Every-Home-in-Australia...html>

Unlike me, Google can put some serious money where their mouth is, and deliver a 1Gbps service:

http://news.cnet.com/8301-1023_3-57481108-93/google-shows-isps-how-to-build-a-superfast-network/

If it is inevitable that we have to spend the money on a new communications infrastructure, then 1Gbps is where we should be. That is what the technology of today can give us. So why does everyone want to spend so much just to end up in the same ballpark we are now?

7. "We will deliver a better NBN."

The criteria for a 'better NBN' are:

- 1Gbps to the home
- at the same cost to the user as ADSL is now
- with a declining user cost over time as technology becomes cheaper, down to \$20 per month after 5 years

Now that would really be something significant for the country and truly deliver the promise of an NBN.

Posted by Steve Waddington at 11:35

Friday, July 27, 2012

Google showing how...

... to deliver 1Gbps to the home in this CNET article

pretty much what I said 3 years ago in a blog post.

But good on you Google, for showing telcos how, I think your view has a much better chance of being recognized than mine.

Posted by Steve Waddington at 13:29

Thursday, July 26, 2012

The Value of a Customer

What is a residential ADSL customer worth? I have heard a few people now mention the 'NBN land grab', where ISP's like Optus, TPG, iiNet etc want to grab as much of the current ADSL market as they can, because on the 'level playing field' of the NBN landscape, where everyone's infrastructure costs are equal, the ONLY differentiator will be retail price, and more customers means better amortization of business overheads.

To understand the value of a customer to a Telco, the place to start is to figure the revenue the Telco will be looking to make - not now - but in the 5 year future, assuming the full realization of the NBN roll out as it stands.

There are three services:

- Telephony
- Internet
- Entertainment; a very broad brush that covers online games, but in terms of consumer dollar means essentially video/TV on demand. A replacement for cable in other words.

Looking at average monthly spends per customer now we get:

Telephony.....\$30
Internet.....\$40
Entertainment...\$90

So \$160 per month ARPU (Average revenue per user) for an ISP that can deliver those three services via NBN.

If we assume a margin after delivery and infrastructure costs of, say 30%, give us close to \$50 per month, or over a 2 year contract, \$1,200.

But why would an ISP expect someone to immediately churn at the end of the contract? The delivery infrastructure is the same, the costs are the same. Why does anyone buy BP rather than Caltex? Only marketing and price as far as I can see. All things being equal, once connected to an ISP of your choice on the NBN, unless there is some compelling reason to change, you are going to stay with that ISP. So maybe the average 'life' of a customer will be 4, 5 or even more years.

Here is the current poll on Slashdot that gives some validity to that guess. The poll responders being mostly in the 'expert user' category, and mostly American.

Slashdot.com user poll 26 July 2012

Telco's, differentiated from non-Telco ISP's, tend to think in 15 year cycles. On that basis, 2 years can be pretty much written off as simply marketing and acquisition cost. Which puts the 'value' to a Telco of a residential customer, that they don't currently have, at \$1,200. At least by my back of a napkin, rough estimate, estimate.

But then think about this; I was told, by a reliable industry source some time ago, that the Big Pond 'win back' campaign of 2 years ago had a reserve of \$900 discretionary allocation they could use to entice customers back to their network. Also, I have heard from several other people, who I have no reason to doubt, that ISP's with their own DSLAMs, like iiNet and TPG, see the value of acquiring a customer at about \$900. And that information is now at least a year old, so maybe my \$1,200 isn't too far off the mark.

Maybe those guys figured out long ago what a residential ADSL customer is worth. I am feeling a little slow today.

Posted by Steve Waddington at 14:14

Thursday, July 12, 2012

Bingo, Bongo

In 2009 we implemented a customer survey program for support issues, when a ticket was closed. It was done for the obvious reasons; to get the feedback direct from the customer on how the case was handled; get the information direct from the source on what we may need to improve on; and track changes/improvements we make over time to make sure they are effective.

The results, graphed over 2 years, showed a steady month on month improvement, and this tracked nicely also with a general decrease in emails to 'complaints@exetel.com.au' and TIO case reduction. Because each survey response is linked to the case ticket, we can see who has worked on the case, and the support supervisor or manager can review any low scores and identify why we didn't make the customer happy.

In our own terminology, we dubbed a high score response - where the average is 4.5 - a 'bingo', and a low score response - where the average is less than 3, a 'bongo'. Bingo's were encouraged, with engineers who achieved consistent high scores being recognized for their work in their quarterly reviews. Bongo's on the other hand were treated very seriously, triggering a 'Negative Customer Feedback' email to be sent to support supervisors, manager and the Directors, and requiring an cause analysis and corrective action to be undertaken.

The directly measured improvement over time, along with improvement in related metrics, showed the program was successful. But still there seemed to be something missing.

By giving such focus to the low score responses the focus was always on the negative, rather than the positive aspects of what was being achieved. And while support engineers were recognized for achieving good score, the recognition was well after the event.

So in February we began the 'Bingo Bonus' scheme, where the support engineer gets a pay bonus for each 'bingo' feedback response they get from a customer; qualified by three simple rules:

- if two engineers are involved in the case, the bonus is split 50/50
- if more than two engineers are involved in the case, that case is not eligible for a bonus
- any bongo response disqualifies that engineer from any bonus for that month

The results speak for themselves. In the December 2011 there were 35 Bingos and 5 Bongos, the average handling score was 1.2 (note on what this means at the end of this post) and the reputation score average was -0.3.

January was similar.

The scheme started mid February, and the early results hadn't changed much - 35 Bingos, 5 Bongos, handling score of 1.16 and an improved reputation score of -0.1

But by May, Bingos had increased to 56, Bongos was still 5, handling score was up to 1.58, and reputation had turned positive to 0.1. Then in June: 75 Bingos, 4 Bongos, handling 1.59 and reputation 0.5. Since we are dealing with averages of hundreds of results, that sort of increase represents an order of magnitude improvement.

The total bingo bonus monthly payout for June is about equal to the salary of 1 level 1 engineer. So to look at it another way, would the addition of 1 more level one engineer equate to an equal improvement in customer satisfaction? Highly unlikely. And it certainly wouldn't have given the other engineers the motivation and incentive to achieve as they have.

Note on what the survey scores mean:

Reputation: Q1 asks what the customer how they rated Exetel before they has the problem, Q7 asks what they rated us after. The reputation score is the difference between the two, averaged over all responses.

Handling: Q2-6 ask the customer to rate aspects of how the case was handled, the result is averaged for each case, subtracted by 3, and then averaged for all responses. We subtract 3 from the score so that when the results are

graphed, the reputation line and the handling line are both in the same scale, and it is easy to see at a glance where the handling line drops below zero. A handling score of 1.5 equals a 4.5 response score average from the customer.

Posted by Steve Waddington at 14:32

Tuesday, July 3, 2012

Exetel Communicatoins (Pvt) LTD Wins International Award

Posted by Steve Waddington at 08:38

Monday, July 2, 2012

EOFY

Well, we ended the year with a bang, both good and bad.

The good parts were:

- Another record sales month for corporate services with over 200 new services signed for the month
- Outbound sales in Colombo firing up, sales targets well on track there and making a significant contribution to our sales figures
- 6 successive months of downward trends for customer complaints
- 6 successive months of upward trends on customer satisfaction surveys
- 4 successive months of improved indicators across the business

Very gratifying to see these results, which certainly have not been achieved without a lot of hard work, some tough, sometimes unpopular, decisions and a few tears along the way.

And the bad part - a literal 'bang!' was the failure of AAPT's IP bandwidth on Saturday:

We are currently waiting for the official RFO, but from an engineering perspective, it is pretty clear to us what happened. It is also clear that AAPT are going to have to take a hard line within their own network engineering group, cut through the finger pointing and blame avoidance that is bound to happen, and fix themselves up pretty damn quick - if they want to retain our, or anyone else's business in the future.

P.S. Edit: Leap second injection has been offered as the root cause.

Posted by Steve Waddington at 11:41

Monday, June 4, 2012

Records Tumble

May set a new record for revenue, and new records in:

- Corporate sales orders, beating the previous sales record of 150
- Provisioned corporate orders, ending on 102
- SMS sales, exceeding 950,000 in a month for the first time
- VoIP minutes
- Residential new business revenue

While residential (ADSL) sales didn't quite match the record month of February, the new, higher value Anzac and Fair Use plans pushed revenues for new orders just over the February mark. All in all it was a great pick-up from April, and we are now at 106% YTD sales target for corporate sales and just a shade over 100% YTD forecast for residential sales.

June, as usual for that month, is shaping up to be another record killer. Unless I have just jinxed it.

Posted by Steve Waddington at 17:18

Wednesday, May 23, 2012

AI Project Culmination

The R&D program we began with the Sri Lanka Institute of Information Technology two years ago is in the final stages of completion. Yesterday we held a press conference, organized by the SLIIT, to announce the results of the project. The press coverage included, all the national TV stations, papers and the local Al Jazera new team.

The presentation I gave follows, but the most impressive part of the conference was Prof. Kumarawadu's demonstration where he was showing question/answer response illustrating the confidence level algorithms that form the core logic of the system. A few minutes into the demonstration of the web based interface, I noticed that the url was <http://calliope.exetel.com.au/> - our live production system. Not some mock-up put together for the press, not some abstract diagrams of an ideal system, but the real, live, in production, result of the development work. Now if that isn't confidence in your work, then I don't know what is.

TV News cameras setting up.

Conference panel, from left to right, me, Professor Kumarawadu - Research Leader, Professor Lalith Gamage - President/CEO of SLIIT, Professor Sam Karunaratne - Chairman of SLIIT, Lahiru Sanarakoon - Research Engineer

Prof. Kumarawadu demonstrates Calliope live to the press.

My presentation:

Distinguished guests, professors, members of the press, Ladies and Gentlemen,

I am the CEO of Exetel and we are a small Australian company and have been operating in Colombo for four years.

From the start we saw the potential for collaboration with academic institutions in Sri Lanka.

There are many difficulties that any company must face starting an operation in a different country. One of the hardest things we found with regard to our academic collaboration program was how to engage in the first place.

I am not sure why it was, but for the first six months when I was trying to establish this program, I could get no traction at all from any of your universities. You would think that the offer of fully funded scholarships and research funding would have some appeal, but it didn't seem that way.

Maybe it is a reflection of the excellent education system and government funding already in place, that private enterprise funding is not needed or wanted. As I said, I don't know.

But, what I can say is that we did overcome that hurdle, and the SLIIT had the initiative to engage with us in, what was to become, three very successful projects.

Before I talk about the project, our AI engine called Calliope, I would like to first speak a little about, what I think, is at least equally beneficial, and I feel a lot of pride for.

The first initiative we undertook with the SLIIT was to provide two fully funded scholarships per year, in the field of Unified Communications. We started that program in 2009, and we now have six SLIIT students receiving 70,000 rupee per semester fully funded scholarships.

But what I am particularly proud of in that project is that the scholarships are completely without encumbrance. That is; we ask for nothing from the students, other than that they maintain a GPA of 3.0. We do offer the opportunity to paid internship work, but it is not mandatory. On graduation, our intention is to offer the students full time work at Exetel. But this also is not a mandatory condition.

Our objective in doing this was not completely altruistic. We wanted to gain access to the 'best and brightest' graduates, who would see our company as the best place to work for them, to begin and their working careers.

We were not looking for indentured servitude. We wanted bright, thinking, capable people, who wanted to work for us. And what I am very pleased to be able to tell you today, is that is just what we have achieved.

We currently employ 13 SLIIT graduates, paying them the highest salaries of any company in Sri Lanka, and providing technical support and network engineering services to our Australian operation that is the equal of any in the world.

Our second initiative has been to offer paid internships to final year students. We have always had great success with this program in Australia, where students have worked as interns, gained full time employment on graduation, and gone on within Exetel to become team leaders, supervisors and managers. In Sri Lanka, three people who started as interns are no full time staff, and we have two current interns.

We have made those investments in good faith. And it is a tribute to the quality of the SLIIT graduates that our faith has been vindicated.

The third initiative is the reason we are here today. The culmination of the funding of of the SLIIT research project into Artificial Intelligence, which we call Calliope.

Calliope is from the ancient Greek, and was the muse of epic poetry.

And bringing this project to fruition has been something of an epic.

So what is Calliope? It is simply this; a system by which a customer use natural English to find out the information they want. One way of doing that is for a customer to call a call centre. And for complex, difficult problems, that is exactly what they should do.

But so many, by far the vast majority of calls to any call centre are simple questions, usually answered somewhere on a company web site or FAQ page. But for the customer, it is far more convenient to just call someone than to try to find the information on a web site, no matter how accessible the company thinks it may have made that information.

So eight years ago we conceived of Calliope as a method to make us a more efficient company, and make the information our customers wanted more available to them.

Our own development stalled in its second year. We had simply reached the limit of the resources a small company like us could commit to such a project. And the concept languished for a few years.

We were not by any means the only company to consider such an approach. There are a number of such AI's in commercial use. But from our point of view they all had the same failings: they were expensive to buy, hard to maintain, and as far as we could see, only able to deliver a valid answer about 50% of the time.

We wanted something better than that.

You may have heard of the recent current 'killer app' for tablets and iPhones called 'Siri'. On an individual basis, Siri is exactly the sort of intelligent agent that Calliope is on a commercial basis. You could say Siri is to Calliope what a spreadsheet is to Oracle database.

Two years ago Exetel began a collaboration with the SLIIT through funding to develop an AI system that was commercially viable. Exetel had three criteria that needed to be met to achieve this;

First the system had to draw from a single source of information - the Exetel Wiki site.

Second, it needed a knowledge base that was easy to maintain. A single source of update information, which was our own company knowledge base - our wiki site, would achieve that

And third it had to integrate fully into our existing customer management systems.

The SLIIT development team have met those objectives to the extent that the systems is and integrated and valuable part for our customer support.

The benefits Exetel and its customers gains from this system are:

- Skilled support engineer time is freed from mundane repetitive tasks that the AI can handle, and can concentrate on

more difficult solutions

- Customers can get immediate responses to all basic questions about support and provisioning, and have direct access to the information they need
- Exetel has reduced its cost for basic support, and increased the quality of support to customers

The successful culmination of this project is one more demonstration that Sri Lanka has much more to offer the world than Tea, gems and garments. For a small company like Exetel to be able to engage and successfully complete such an endeavor with the SLIIT stands as an example of the intellectual investment value Sri Lanka can offer.

And so in closing I offer my congratulations and sincere thanks to Professor Kumarawadu and his development team, and the SLIIT for engaging in this project with us, and to make our vision and initiative a success.

Thank you.

Posted by Steve Waddington at 13:28

Tuesday, April 24. 2012

It's Not Fair

We extended the offer for Fair Use plans through April, and will most likely keep them available for May as well. They have proved pretty popular, though, as the only 'no contract' plan on the market, and better value over 12 months than any 'unlimited*' plan, sales have not been as high as I would have thought. Still, it was a good enough deal that I bought it myself for my family use - so no more excess use surprises for me when my son at Uni downloads his semester video assignments in RAW format.

I guess I can understand why some people might be reluctant to buy a Fair Use plan - where the clearest definition of Fair Use we can use is based on an average of use. And if that were the only criteria, maybe we could take an historical average and publish that as an indication. Well, not really - so I am told by our compliance officer - if we publish a figure we have to stick by it. The other thing is, the average use is quite low, less than 100GB overall, and that is just not an appealing number, where legitimate use nowadays can see in three or four times that downloaded in a month. So 'Fair Use' does require some two way trust; the ISP needs to 'trust' the user not to download the entire Internet, and the user needs to 'trust' the ISP not to penalize them for some arbitrarily low, albeit average, download amount.

Bottom line, no matter how good the deal, Fair Use plans are not for untrusting people. (and I am not using the term 'untrusting' as a criticism of anyone, maybe 'savvy consumer' or 'healthy cynicism' would be better terms)

So what to do?

We came up with the 'Anzac Special', no contract, 1TB allowance, \$55pm - same as the Fair Use plan, but for people who want certainty in quota rather than certainty in excess usage charges.

The plan should be on the web site later today, and we will see how it runs to the end of April and decide if we want to extend it into May.

Posted by Steve Waddington at 12:32

Monday, April 2, 2012

1300 Days on the Spanish Main

Nearly four years ago we put in place a system to forward on copyright infringement notices that identified an IP address to the account to which that IP address was assigned. The system has been running faithfully for, today, 1299 days, and sending a report each day. Graphing the results shows an interesting trend:

Y axis is notices per day, X axis is days

Considering the amount of bandwidth available to residential subscribers has almost quadrupled over that time, the possibilities are:

- Copyright infringement notices and anti-piracy campaigns are effective over time
- P2P stealth has got much better
- Users pirating have got pretty much all there is to get, and there isn't much left
- Copyright enforcement companies have become less active

I don't think the last point is true, their systems would be pretty much automated by now I would think, so there is no particular reason so think they have scaled back.

P2P stealth I don't think is the reason either. No matter how much encryption, both ends still need to know the IP address and the file name, and that can't be hidden transferring a file to/from a 'honey pot' peer.

So, is it Hurrah! for the good guys or has all the gold that can be plundered been plundered from those Spanish galleons?

Posted by Steve Waddington at 10:16

Friday, March 23. 2012

A Sudden Brain Storm

We were talking about 'unlimited' plans on offer in the market, and how they all have one or more asterisk, tilda or some other super text that points to a 'fair use' policy that explains how 'unlimited' doesn't really mean unlimited. And why the ACCC. for some reason, requires 'unlimited*' actually be 'unlimited'.

So, why not, we mused, just call the plan 'Fair Use' up front - flat rate, no metering, no shaping, except in conditions where it is required - which is what the fair use policy says. Just instead of 'hiding' it and hoping no one notices the tricky conditions of 'unlimited*', have fair use up front, and explain what makes it unlimited. In fact, do away with any word as possibly misleading as 'unlimited' and state plainly what is being offered.

Maybe there is something we all missed in this, why no one does it? I guess we will find out.

The 'soft release' of our Fair Use \$55 and Fair Use 65 plans will be on the web site later today. It will be a limited time offer, and we will see how it goes.

Posted by Steve Waddington at 22:00

Tuesday, February 14. 2012

A kinder, gentler NBN?

I see Delimiter have noticed our new NBN pricing. A fairly well balanced article, but we do have a high use plan which they didn't mention - for a 300Gbps quota. And, famous last words, 300Gb should be enough for anyone, right?

Another important difference is the Exetel quota is download only, in theory, that makes it double the quota of other ISP's who include uploads in the quota. Though in practice, it works out to 15-20% extra, or whatever the ratio of downloads to uploads for a particular user.

I don't know why it is that everyone reporting on the NBN, Delimiter's article included, seems to be using this early 'get it off the ground, feel out the market' pricing as proof of future NBN pricing. It can't have escaped everyones attention that introductory pricing is often much cheaper than the price something will be once it had gained market acceptance. I don't think the NBN will be any different.

Knowing the NBN cost, the number of people who will use it, and the way it will be paid for, it can't be too hard to work out what the price will be for access. Something that hasn't escaped the Coalition's notice, but does seem to escape journalists. Oh well.

Posted by Steve Waddington at 10:37

Wednesday, February 8, 2012

SMB

We have been running a trial marketing program for a bundled SMB product for the last month. I really like the idea of this sort of product, it meets a 'niche' - but that is a very big niche, of over 1,000,000 businesses.

The marketing is expanding throughout February, and we will be moving support engineers into a specialist SMB support group in March. Which is hand-in-hand with the plan we started in November to increase the number of sales and support engineers for both residential and corporate - ahead of the 2012 sales forecast.

Well, the sales numbers are right on track (better than forecast for residential in fact), so our hiring plan has been right on time.

It is very nice when plans coincide with targets so well, even nicer when the timing works out so well too.

Posted by Steve Waddington at 11:25

Monday, February 6, 2012

NZ Expansion

We have had a POP and bandwidth to New Zealand for some time now, but have not had particularly good local loop 'reach'.

That is likely to change soon, as we re-assess the potential, and complete an agreement with a wholesale supplier for local loop circuits. And also significantly increase our bandwidth there.

The price for a cross-Tasman circuit, Glenn tells me, will be very competitive. But my first reaction was 'That much!', because, for some reason, bandwidth to NZ is very expensive. I just don't know why that would be - apart from opportunistic carrier gouging. So I guess we will have to see if we can still offer a service at a price people will want to pay.

Posted by Steve Waddington at 15:48

Friday, February 3, 2012

A Footnote

January was a record sales month. Yesterday we had one of our highest order days ever.

John and I spoke about why that was briefly last week, and neither of us had a good explanation. My 'throw away' line comment was that it like we had won some sort of battle.

Maybe when you have had your head down plowing through heavy clay for so long, there eventually comes a time when you hit some soft loam, and you don't even realize it at first.

I think I have strained enough metaphors to breaking for today.

Posted by Steve Waddington at 08:37

Thursday, February 2, 2012

Farewell John

I have worked with John for the last 16 years, at four different companies, and been his business partner for the last eight years.

We started Exetel with the objective of creating a 'perfect company'. We faced many challenges, some that would have undoubtedly overwhelmed anyone with less indomitable determination than John.

He was the toughest person I think I will ever meet. The most honest person I have ever met, and one of the kindest. But above all he was unique, with the clarity of vision and sharpness of mind that was simply awesome.

He has been my manager, my mentor and my friend.

Farewell John.

Posted by Steve Waddington at 09:23

Friday, December 16. 2011

Representing Exetel's Network II

This is the 3rd iteration of the high level diagram we will use in 2012.

There will no doubt be a few more tweaks before it sees commercial release, but I think it captures the essence of what we want to convey.

'Underneath' this diagram will be more detailed, but still conceptual, location and service specific diagrams, about 10 in all, that will show how services are supplied. And finally, there will be the engineering design diagram, individual to each customer, that is attached to a proposal, showing in detail what is being supplied.

Posted by Steve Waddington at 12:35

Tuesday, December 13. 2011

Public record of What you have downloaded

Here is an interesting site.

Should give a few people pause for thought.

Posted by Steve Waddington at 11:36

Friday, December 9, 2011

Representing Exetel's Network

I have been thinking about how to redraw the Exetel network for sales and marketing material. The network is long past the stage where it can be represented in full, even taking treating each POP as just a circle and only putting in the WAN links.

The challenge is to show 'at a glance' the main features of a network with multi, multi GE and 10GE interconnects, dozens of peers, data centre locations and services, that spans three countries. At the same time contain no false statements (like using a port speed or supplier network capacity to represent a circuit speed), and be accurate; as a summary of a network not a 'lossy' compression.

Yet even if a diagram did that, it's not the whole story of the Exetel network. There is one massively important part of any network that rarely, if ever, makes it onto the diagram.

After a lot of thought, this is what I came up with:

In this iteration I haven't put in the specific WAN circuit sizes, for two reasons; One, given our no contention policy, the size of any one circuit doesn't matter, as long as it exceeds the peak bandwidth at all times. Two, Due also to the no contention policy, circuits are increased often - faster than the print/distribution cycle of a brochure. And, in any event, anyone interested in the fine detail of a location or data path can ask, and be given the up-to-the-minute value at that point.

Only a draft at the moment. It will be interesting to see how the bright and capable minds in sales improve it by the time it gets to 'brochure' stage.

Posted by Steve Waddington at 16:35

Wednesday, December 7, 2011

Workplace Safety

No one can deny that walking around on a big flat field during a lightning storm is a dangerous thing to do. Quite rightly, workplace safety should be observed.

No one can deny that an aeroplane is made of conductive metal, and much, much taller than a person.

So all those planes sitting on the Perth tarmac yesterday were being constantly hit by lightning? It must have been terrifying for the passengers sitting on the planes.

The Perth terminal, last time I looked, has plenty of steel and metal superstructure, and is taller again than the largest plane that docks there. And I would be very surprised if it didn't have substantial lightning rods higher than that as well. It must have been deafening as bolt after bolt hit the building.

With so much metal sticking up in the air, the safest place to be would surely be on the tarmac apron between planes on one side and the terminal on the other, being effectively surrounded by a big Faraday cage.

But physics is clearly only one factor when it comes to workplace safety. Because on the other hand, sitting inside playing euchre all day because it was a bit wet out is probably a better choice.

Posted by Steve Waddington at 12:00

Monday, October 31. 2011

The Price of Loyalty

Over a decade ago I worked for a wholesale IP supplier. It was only a small company, but, because of some innovation and foresight by one of the founders, ended up supplying about 30% of the non-Telstra bandwidth used in Australia at the time.

The great innovation that they pioneered was the use of uni-directional satellite transmission as an alternative to the, then, much higher cost of terrestrial bandwidth. As I recall, satellite bandwidth (on the old, wobbly PAS2) worked out around \$5,000 per Mbps per month, compared to terrestrial (pre SCCS) at around \$17,000 per Mbps per month.

With a delivery cost less than one third of the alternative, many ISP's found they could overlook the additional 200ms latency, the bi-annual sun transit outage, and the more frequent downlink signal loss (for any number of reasons). So despite the less than perfect service availability, the company had many loyal customers.

Until there was some competition, in the form of much lower cost terrestrial bandwidth. And the 'loyalty' evaporated in a direct relationship to pricing as the cable price approached parity with satellite.

I am not drawing any negative conclusions about ISP's making the only decision possible for them - which is to take the lower cost, better performing, more reliable product. It would have been stupid for them not to. Rather, it's just the fact that many people would start of justifying their decision by saying 'I have been your loyal customer for many years, but...'

Which made me wonder, how is that loyalty?

Loyalty would be where the customer just kept paying whatever was being charged, in the face of all good business sense, forever. Or a company 'loyal to its customers' forgives any non payment for services. Or a patriot, loyal to their country/ideal/football team, lays down their life for it, or at least continues to support it when they are at the bottom of the ladder.

That is loyalty, and that is its price.

So what is the real, commercial price of 'loyalty'? In my experience it can be defined as exactly one cent.

For example, many people I knew in my home town of Albany would go 'loyally' to the petrol station owned by their mate/relative. Until, the price was one cent a litre cheaper somewhere else. Then they would have no qualms about driving to the other side of town (only a 3 minute drive, but still) to fill up.

Or, why are prices \$x.99 and not \$(x+1).00?

Going back to the wholesale bandwidth provider I was talking about. Customers would remain loyal until the price per Mbyte was 1c cheaper somewhere else.

Still, I suppose 'loyal' is less than a mouthful to say than 'I have been your opportunistic customer, while your price was the cheapest and while it suited me'. Not that there is anything wrong with that. That's what we all, as customers, do. Maybe it is even hardwired into our brains - to constantly weight the price/quality/value, and feel ill at ease if we do not have what we consider the optimal mix.

Interesting that of all the ISP's around in 1995, and still in business now, the most successful by far is the one that (as told to me by the CEO at the time) aimed to always be the lowest cost producer.

Posted by Steve Waddington at 12:37

Thursday, October 27, 2011

NAT; Poor Mans Firewall or End User Savior?

I was intrigued by this article, which seems to be advocating the benefit of IPv6 is to allow direct exposure of user (student) PC's to the Internet. An interesting view. I wonder if they have thought about that at all.

I can't recall the source, it was a few years ago now, but it was a reputable authority, and I had no reason to doubt it was true; that a new windows PC connected directly to the Internet, would be compromised and incorporated into a bot-net in less than the 20 minutes it would take to download the security patches from Microsoft.

My own view is that, without NAT, the burden placed on ISP customer support just to say 'Not our problem, you are responsible for the security of your computer' to customers that had been hacked by bot-net collectors would more than double the cost of Internet services. Assuming an ISP could even get away with saying that, which they couldn't, and some officious govt watchdog would no doubt see an opportunity to justify their meaningless existence and step in and make it mandatory for the ISP to fix that too.

Effectively, without NAT, ubiquitous residential Internet would not exist.

IPv6 solves the depletion problem of IPv4. And that's it.

Still, a great opportunity to make lost consulting money either way. It looks to me like the Y2K of this decade.

[Aside: remember the panicked articles earlier in the year about 'The last IPv4 address will soon be allocated'. While they didn't actually say the Internet would end on that day, none of them went out of their way to point out that it wouldn't either. Well, it hasn't, has it. And shows no sign of doing so. But who gets column inches or sells copy saying 'IPv4 nearly gone, but not really a problem for a few years yet. Experts say'.]

Posted by Steve Waddington at 17:20

Wednesday, October 26, 2011

One way to end Copyright Theft

Speaking of civic virtue and social responsibility, I was in Tokyo a couple of weeks ago attending NTT's annual conference. One of the presentations was talking about bandwidth use - predicting the, now unsurprising, exponential growth. But what was surprising was a completely unexpected (by me) sharp downturn in one area of bandwidth use, namely; P2P.

The graph shown was a measurement of OCN bandwidth since 2004, showing a consistent average of a bit over 40% growth per year for 'down' traffic, and the same growth for 'up' traffic, until 2009 where traffic dropped from 200Gbps to 140Gbps over the next two years.

The major contributor to the increase in down traffic was explained as the increasing use of streaming content. But why the puzzling decrease, against all trends, of up traffic?

Quite frankly, I found the explanation astounding.

Apparently, in 2009, the Japanese government issues a 'stern warning' to the nation about the security issues of P2P, and that is was not socially acceptable to download unauthorized copyright content. And so, people stopped doing that.

Draw your own conclusion about how strong the moral fibre of a society that can do that must be. My thought at the time was 'if only'.

Posted by Steve Waddington at 09:46

Wednesday, July 27, 2011

Internet Speed

Here is an interesting article from Akamai.

If average peak speeds are, as reported, 14.4Mbps, which takes into account people still on 1.5Mbps and lower ADSL circuits, as well as people on ADSL1+ and ADSL2 circuits, then by my rough calculation, it puts the average ADSL2 connected user on a peak download speed of around 16Mbps.

Compare that to the connection speeds published by NBN, with the entry level of 12Mbps.

What is the government spending all that money for again?

But what was the 16Mbps peak in 2009? Well, it's a long time ago now, but my guess is the take up of ADSL2 - the carrier ADSL2 network backhauls had less contention.

Any service provider using the NBN will be able to exercise control over peak use by the amount of bandwidth they order to connect to an NBN POI. The 'level playing field' price for that bandwidth, represents between a 50% and 70%, or more, discount on current carrier interconnect pricing (obviously depending on the carrier, and the company buying it).

Which would have to mean that any ISP now following a policy of zero contention could be safely assumed to have no real problem continuing that policy based on the NBN price model. And I can only agree with our Sales Director Glenn; broadband users are unlikely to increase their spend on service 'just because' it's the NBN. Rather, they are going to replace an ADSL service with an NBN service, IFF the change over cost barrier is low. In other words, someone spending \$35 or \$45 per month now will take an NBN service at that price, but someone spending \$35 today is not going to suddenly say 'hey, it's the NBN, I think I will now spend \$70 per month on Internet access'.

That is going to put most users on either a 12Mbps or 25Mbps circuit speed plan type - for Exetel customers - or less, if you take the industry average of published prices.

I think that is going to work out at an average peak speed of 14-16Mbps. But maybe someone with a better grasp of stats can correct me.

Posted by Steve Waddington at 16:09

Monday, May 2, 2011

1,000 Baht

On my way to Sri Lanka this weekend, I had an overnight stop-over in Bangkok. What I noticed this time more than others is just how good the Thai people are at the art of the 'up-sell'.

It was amazing, I don't know why I hadn't noticed it before. Just about everything that starts off with a price of 250 or 350 Baht ended up with me parting with the better part of 1,000 Baht. And being very happy to do it.

Here's one example; With a couple of hours to kill between hotel check out time and getting a taxi to the airport, I decided to treat myself to one of the great treasures of that country - a Thai foot massage.

250 Baht for a one hour massage, or about \$7 at today's exchange rate - very reasonable.

'Sir, would you like a foot manicure as well? We have lady who very good. Only 250 Baht'

Of course I would, why not, since I'm here.

Then, feeling totally relaxed halfway through the foot massage/manicure 'Sir, your fingernails long, would you like hand massage and manicure? 150 Baht.'

What's that. four dollars more? Sure, please, that would be nice.

'Sir, you like some herbal tea? 70 Baht'

Sounds good, yes, thank you.

'Sir, you like fragrant oil neck massage to finish? 100 baht'

Yes, lovely, thank you.

So there you go, 250 Baht going in is 820 Baht, plus tip, 1,000 Baht going out. But the important thing is, I was really happy to spend the money. The masseuses were friendly, smiling and pleasantly chatting away to each other the whole time. And I came out feeling that merely spending four times more than I expected was an absolute bargain.

Thais - the masters of upselling, because once you find out how good the value for service is, you are always happy to spend more, to the extent you appreciate them suggesting how you can do that.

It made me think - I think most of our customers are pretty happy with the money they spend and the service they get (obviously, because there are very few barriers for those that aren't to move somewhere else). But just how much better could we be?

Posted by Steve Waddington at 13:02

Friday, March 18. 2011

PeerApp Ultraband Cache, 3+ years on

Three and a half years ago we deployed our first data cache. In that time the software has been upgraded many times, and 'tweaked' for performance many more. The hardware has been upgraded each year. And since April 2010 two more cache servers were installed, one each in Melbourne and Brisbane.

The milestones, at least for the last 12 months, can be nicely tracked on the MRTG graph:

So what can I say about it after 3 years? There have been pro's and con's of the system along the way, but at every point it has proved better to have than to not have when all things are taken into consideration.

The benefit of hindsight (and record keeping) shows the one and only performance metric of any long term investment - the net dollar benefit over time. Simply put:

Cost

- + a little more in engineers and network administrators time
- + some early port redirection issues we would have preferred to have done without

Saving

- 50% of the cost of alternative bandwidth

I can't put any figure on the two costs other than 'something', but we got through it, and didn't ditch the system, so it couldn't have been too bad or too much.

The cost savings though, I can calculate, and it comes out to \$499,500 (net, after all costs) over the three years. Not a huge saving, but never the less, quite significant given our very, very competitive market and ultra-slim margins.

Posted by Steve Waddington at 15:19

Thursday, March 17, 2011

MPLS - Sold by parrots, bought by the weak minded

Not that there is anything wrong with MPLS either inherently or by the fine vendors whose products support it.

But just why do people, meaning the people making the purchasing decisions, buy it for their company?

Here is a valid reason; a large, national network, where the network owner needs the dynamic, self managed, flexibility of 'control over the entire network routing' that MPLS provides.

In practical terms this means; the Service Provider gives you a BGP session from your CE to their PE router for each location, and you, the network owner/customer, do the rest.

Sound like you? Or do you have a network where the service provider gives you a 'fully managed service'? Or perhaps 'A fully managed MPLS service'?

If you are almost every business in Australia, you are in the second, above category.

So now all you have to ask yourself is 'what difference does the MPLS in the description of the fully managed service I am buying make?'

Because the packets coming out of the CPE to your local networks really know the difference between MPLS or any other method of delivery, don't they? Yes, they are so much more relaxed and happy, and their payload of data is just that much better for it, isn't it?

Excuse my extreme cynicism. It just sickens me to continually hear the snake oil pedaled by the charlatans of the industry in this country, that is bought up by their ever eager audience of non-competent decision makers.

If your packet latency is 20ms between two points, do you really think it makes any difference if it gets there by MPLS, GRE, IPSEC or just plain IP, or any other TLA, FLA or nLA someone might dream up in the future? Of course not. A packet is a packet is a packet. Your data could just as easily be encapsulated on a RAM stick put in an envelope and posted to you. You might notice the latency difference in delivery, but your data sure as hell wouldn't know any difference.

Here's a hint; if you see MPLS anywhere in the proposal your service provider gives you, and you don't know what a CE or PE is, then the only thing MPLS means to you, is you will be paying somewhere between 200% and 500% more than you need to for your circuits. But don't worry, your CFO might not find out your blunder this year, or maybe even next year - and who can think further than that into the future.

Oh, incidentally, do Exetel use/sell MPLS? Of course we do. Our routers and network uses the same sort of equipment and OS's as everyone else's. There are some very good reasons for us to use MPLS in our core and connecting to some of our suppliers. It is also readily available for our customers. But, I have to tell you, to date, for every one who has asked for MPLS (always on the basis of 'their last supplier did it, or 'we must have it/no, I don't know why, we just must'), we have always found a more efficient, cost effective way to deliver the same services without MPLS.

Posted by Steve Waddington at 10:41

Wednesday, March 16, 2011

The year of Now

To: Exetel Staff

Ladies and Gentlemen,

I have mentioned it to a few of you, and I think it only fair to tell the rest of you first hand.

My New Years 'resolution' was that 2011 would be the year on 'now' as far as Exetel is concerned.

Over the years of our operation, I have seen one of the founding principles on which we based the company slowly, slowly, slip away. So slowly in fact that I realized, eventually, after thinking back, that even I had become complacent with the downhill slide.

So, that is what 'new years' are good for - you can make a decision that, whatever happened in the past that you weren't happy with, you can take action now to stop it making you unhappy in the future.

And so I have.

The principle to which I refer in this case is the immediacy in which things were always done. That is; every email or request to anyone elicited an immediate response; any request for a change had an immediate response of:

- yes, it can be done, and I am doing it now, or,
- no, it can't be done as you ask, and here is what I suggest would achieve what you want, or
- in the rare cases it is warranted - I am not sure I fully understand, can we please discuss it (and then I will do it right away).

Everyone in the modern world talks about reduction 'stress' in the workplace, and a better 'quality of life'. I can tell you with the absolute certainty of long experience and careful observation over many years, that once you have a place to work that stops the rain coming in, and an income that pays your essential bills, the ONLY thing left in your working life that creates any stress at all is your dependence on your co-workers to respond reasonably to your requests.

The Directors of Exetel do their very best to address the first two of those factors - as much as any business of our small size in a very competitive market can do.

Only each of you, making a decision in your own mind, can do anything to address the one, remaining, very large factor.

Think about it:

- What do we hate most about our suppliers? They take ages to respond and only sometimes meet the dates they have set.
- What do we despise about our competition? We are told time and again by customers who have moved to us that we provide fast, accurate responses to questions, that they never got from their previous supplier
- What do we loath about our contemporaries in the Industry? Their acceptance of slack practices in their own organizations means it takes is days or weeks to get any information we have asked for.

So you can see the common thread. As soon as one person accepts that 'by the end of the week' is an ok time frame, then the person they need to give the answer too has to be told 'by the end of the month', which practically equates to; nothing you want to get done can be done in less than a year, but probably more.

We hate those qualities in others, we should never, ever, accept them in ourselves.

There is only one time frame to do something, and that time is now. This is the year of now.

Regards,

Steve

ps. If you are a manager or supervisor, it is part of your job to preempt what might be asked of you. It is 100% not your job to be asked to deliver something, and then suddenly have to have a lot of meetings and take weeks or months to come up with what you have been asked for. Sometimes, even the best of us will be caught flat footed, but that is when your qualities as a leader will shine through, and you will make a response based on the facts you have to hand right now, and the best of your ability.

Posted by Steve Waddington at 16:36

Wednesday, November 10, 2010

NBN Questions

I was ask to reply to some emailed questions from a freelance journalist the other day, who asked:

Q. > I wanted to look at it from the perspective of ISPs, who will be primarily responsible for providing these types of services, and understanding the possible commercial return.

>

> - Are you considering offering services such as IPTV, videoconferencing package, or e-health services? What would this deliver to customers and how much would this cost?

A. We have offered IPTV in the past, and IPTV is currently actively offered either free or at a small cost by some ISP's now. When offered as a subscription service, even at a very small (\$2 per month) fee, it did not prove very popular. From a base of 50,000 potential users, the take up rate was less than 0.2%.

To be of practical use, IPTV requires a set top box, or a computer dedicated to that purpose, or a method of retransmitting data from a PC to a TV that is not affected by any other use the PC is put to (and that the PC is powered on all the time). The latter two methods tend to be 'too much bother' for a user not otherwise interested in the technology, while a set top box adds extra cost and places the IPTV service squarely in competition with satellite/cable TV and TIVO type services - which IPTV has a great deal of difficulty matching for interesting content at a similar price.

There are a large number of commercial video conferencing product available on the market now. So too are there very many personal products available - as can be seen using any chat service (MSN, Google etc).

Since the NBN is not being built as a commercial grade service, it is unlikely to be suitable for commercial video conferencing purposes (for example, the shared transmission method will not allow the QoS controls needed for a commercial grade service, and possibly there will be security issues). Which in any event are available over, and widely used on, metro Ethernet and SHDSL type services. We have many customers, and we ourselves use this in our business. The cost is much more a factor of the level of picture quality the user wants, and the amount they are prepared to spend on equipment.

As a 'home user' orientated service, the video conferencing therefore falls into the private use category - which is addressed very well by current broadband products in the \$20-\$50 range. (you may even recall before broadband, millions of people still used video chat quite well on 56k dial-up).

I am not sure how you are defining e-health. But in general, the same is going to apply as for video conferencing. I would think privacy and QoS issues would prevent the NBN from being used for doctor-to-doctor or doctor-to-hospital consultations. I suppose if every household also installed \$20,000 of commercial grade high resolution audio visual equipment, then it would be possible to use the NBN for doctor-to-patient consultation - except were tactile examination is needed. But that could also be done with ADSL services as they exist now.

Q. > - What is the market demand for these services now and how will this look in eight years time when it is proposed that the NBN will deliver 100mbps internet connections?

A. Our 'toe in the water' site in Tasmania indicated that when offered for free, the demand for NBN is less than 1% (in balance, the sample size is small and the demographic may not be indicative of other areas). Never the less, it is a matter of fact that all households in the NBN footprint already have access to bandwidth in excess of 100Mbps for broadcast services via cable, satellite or free to air TV. I think in eight years time the Internet, whatever the media used, will still be an expensive way to deliver those services.

However, it would be a total Luddite that didn't think that over a decade 'something' wont come along that will require a) much more bandwidth, b) requires real time bi-direction data, and c) is a killer app everyone will want.

But that has to be balanced against:

- improved compression methods and greater processing power, requiring less bandwidth
- improvements in electronics, providing greater bandwidth from existing copper and wireless

- increase of mobility devices (eg iPad, iBook etc) and greater requirement for portability

Which will make a fixed line product less relevant to the user base, and comparatively more expensive

Q. > - Is it commercially viable to offer these services now? If not, what market conditions are required for this to happen?

A. People are using all services the NBN offers now on ADSL and 3G/HSPA. Any ISP that makes their financial information available will show if they are commercially viable or not.

Q. > I'm keen to hear your thoughts on this topic and how it can be explored.

A. The NBN, in my view, is essentially a mechanism to remove the Telstra monopoly of copper to every house. Maybe it is worthwhile for that reason alone. Or maybe what could be explored is a more cost effective way to do that.

Posted by Steve Waddington at 12:07

Monday, October 11. 2010

The Reverse Terminator

Review time once again; that all too important exercise of, formally, asking each person within the 'umbrella' of your responsibility how they are going, how would they like to be going better, and how long they think they will keep going for. Because, what else does anyone really need to know other than that?

The 'traditional' review questions along the lines of 'targets/performance' are discretely measured along the way. Certainly in terms of any engineering/systems function, it is very hard to argue with the minute-by-minute data from the monitoring system. Nor can the goal of 'zero faults' be argued against the actual achievement of the monitoring reports that everyone has the same access to.

They change a little from time to time, but these are the review questions for this quarter:

1. Were your job goals met for the previous quarter?
2. What can be done to improve your current duties?
3. What would you like the me/the directors/anyone to do to help you with your career?
4. How long will you stay with Exetel?
5. What do you like most about your employment with Exetel?
6. What do you like least about your employment with Exetel?
7. What are your job goals for the current quarter?
8. Anything else you would like to discuss

So, six questions to find out the three most important things. One question to checkpoint that the 'automatic' metrics are working, and one question to catch anything else anyone feels is important. I have no idea how that agenda would fit in with an HR Managers view of how things should be done. But then, I have been to HR review meetings, and the only thing I ever got out of it was the impression the HR people found their job pretty depressing, and we were both glad when the mandated 90 minute meeting time was up.

The review questions seem straightforward to me - which is exactly how they are intended to be. Never the less, I have never held a review where I haven't been surprised by responses to one or more of the questions - and that is always where the most value of the exercise comes from.

Speaking of exercise, I took a couple of days leave during the school holidays and went camping in the Stirling Ranges with my two younger sons. For a continent as geologically stable as Australia is meant to be, I was surprised to find how much steeper and higher those mountains have become in the 25 odd years since I last scaled them.

Our campsite was to the west of Bluff Knoll, and in the very early morning the valleys are in darkness while the peaks literally glow from the rising sun. It only lasts about 15 minutes, but we could actually watch the terminator sweeping the wrong way from west to east as the sun rose over the eastern mountains behind us and the shadows retreated down the valleys to in the west.

Posted by Steve Waddington at 11:45

Wednesday, September 29, 2010

Fun and Games with the NBN

We have our very first users connected. Qdos to Paul and his team in provisioning, and Brendon and his team in F&A for their truly heroic efforts of perseverance in the face of, at times, seemingly overwhelming bureaucracy.

I can hardly wait now until my suburb is connected - only two years at most (or whatever the latest national roll out plan is).

I was having a think about all the things someone will be able to do with the massive bandwidth available through the NBN, and comparing that to what can be done now. This is what I can up with:

I don't know, maybe I have just become way too cynical. All those 'field of dreams' networks that were built in the 98/99 era got used, eventually, with only two or three successive owner bankruptcies along the way.

It's late, and I'm tired. Maybe someone else can come up something it would actually be useful for.

(yes, I meant to put Qdos. We all remember QDOS from Seattle and later Sinclare, right?)

Posted by Steve Waddington at 02:49

Thursday, September 16, 2010

It's like deja-vu, all over again

(Yogi Berra)

I refer to this latest effort from our nations largest ISP and Telco.

I do recall, was it three or four years ago, when a similar thing happened. My memory is a little hazy on the detail, but I think it was a two week outage for some customers, though it may have been less. Which was due to somehow the email servers becoming unable to cope with the load any more and all crashing in some unrecoverable way. The solution, as announced by their PR person, was to ship in new servers by the truckload. The hardware vendor account manager must have thought all his Christmases had come at once on that day.

Four days only this time though. I wonder what the problem could have been (or still is), and what the fix is? It's probably 90% software and the other half hardware related.

Still, even Napoleon had his Watergate.

And as their spokesman says the Telco/ISP "... apologises for the inconvenience and it aims to update the service status page soon"

Well, as long as the service status page is updated, that's alright then. It's almost as good as actually fixing the problem.

It's true what Yogi Berra says, A nickel ain't worth a dime any more.

Posted by Steve Waddington at 15:45

Wednesday, September 15, 2010

480Mbps Speed - Not True!

I can't understand it. The specification is quite clear. With my USB 2 port, I should be able to get 480Mbps download from the attached hard drive. But I can't. It's just not good enough. I am not getting what I paid for. If only there were a TIO for USB. Why doesn't the government do something about this outrage!?

Ridiculous of course. We all know there are many other factors that affect a data transfer rate than the raw engineering specification of an electrical circuit. Bus speed, disk I/O rate, etc etc.

I predict though, that we are going to see claims like that more and more in press articles relating to high speed circuits from NBN and its competitor networks, if this CNet article is any indication.

I mean, come on now, did we not know geographic distance was a factor in data throughput speed? We are talking about TCP right, syn/ack takes a finite time. Do the math for a 64k window size and, even though it involves the advanced mathematical concept of multiplication, 2.5Mbps is about the best data rate you can expect for a single stream flow on the other side of the world.

The article is also pretty devoid of a number of other things that might need to be known to form any sort of realistic opinion. Like what is the connectivity of the test sites used for example? Looking at that speed test site, none of the sites stand out as locations with giga-circuit bandwidth. What if they have only 100Mbps, or less, connectivity themselves? Probably they don't, but we have to take that as an assumption of which we have/are given no data.

Let's say that (as could well be the case, though we don't actually know) the test site server are located in some well connected data centre with a Gig Ethernet connection to the data centre backbone, and, are capable themselves of multi-100Mbps sustained throughput. Then we just have to extend our every growing string of assumptions just a little bit, and assume they are doing nothing else but being sitting there waiting for the single CNet tester to exclusively run their test. because if they were doing something else, or someone else just happened to be running a test on them at the same time, then we would have to factor in the impact that might make the results. Far too hard to do, much easier to just make the assumption that they aren't.

Anyway, I have now reset my computer and hard drive, swapped the cable, and I still can't get more than 80Mbps download from my USB drive. I am taking it back to the shop to demand a refund.

Posted by Steve Waddington at 12:24

Tuesday, August 17, 2010

The end (of IPv4) is Nigh-er

One of the sessions at the next APNIC conference will discuss the last allocations of IPv4. Still a little ways of yet, but as the abstract says 'we move into the final stages of IPv4 address distribution'.

I like to think our own use of IPv4 has been fairly tightly controlled. Six years ago we applied a considerable amount of discipline to, what seemed like then, the massive amount of address space APNIC granted us on our first application - a whole /17. Since then we have added two more /16 blocks, and continue to manage our subnet use down to each /30. From memory (it has been a while since I have read the APNIC allocation policy), APNIC requires demonstrable use if 80% of currently allocated space before further applications will be considered.

For our part, we have assigned large blocks of each new allocation of address space to each POP, based on our estimate at the time of comparative use. But when push comes to shove, we don't hesitate to re-allocate down to whatever subnet level is necessary (as the APNIC policy requires in any event) to give us as close to 100% use of the address space as we can get.

It makes for some messy route tables - but nothing that the IGP can't cope with. That is what they are for, after all. The only downside, if it can be really called one, is sometimes not instantly recognizing the POP location of an IP address. Which is what traceroute is for, isn't it?

With our metro ethernet and SHDSL sales over the last 18 months, we have been burning through /28, /29 and /30 subnets at a fair old rate. Whereas in previous years we would find ourselves using around 8-10 /24's a month, we are now averaging around 14.

We do have a spanking new Ipv6 allocation, all ready to go... As soon as anyone else at all (ie our suppliers) can accept it. Right now though, it is about as much use as a fax machine in 1843. In the mean time, I guess we will just have to make our contribution to APNIC's last /8 assignment and apply for some more IPv4.

Posted by Steve Waddington at 10:41

Monday, August 16. 2010

Brisbane POP Upgrade

We have a major build taking place at the Brisbane POP this week which will involve, among other things, upgrading the mid range LNS's we have there with a new Cisco 10000.

On the network engineering side, the work is pretty straightforward. Essentially the new router is a drop in replacement requiring very little change in code and simple unplug and replug of the various connecting circuits from the old routers to the new router.

In practice though, nothing is ever that simple. The rack the work will take place in is close to full, and before the main work can even start, it will need to be reorganized and re-patched to make room for the, much bigger footprint, new router.

So the work will take place over two days - or rather two nights, with reorganization taking place in the wee small hours of Thursday morning, and the router replacement taking place on Friday morning. In the 'best case' scenario, we are not expecting any outage during the reorganization phase. Then 2 x 15 minute outages when the routers are swapped over.

For the Thursday morning work, I get a Sod's Law factor of 12.22%, and for Friday morning, 13.16% - the relatively low percentages coming mainly from the skill of the engineers who will be doing the work.

Posted by Steve Waddington at 12:05

Wednesday, August 11, 2010

I'm not a tech head....

Oh wait, I am a tech head. But I don't think you need to be one to see the bleeding obvious - \$43bn and no plan vs \$6bn and no plan... which do I choose?

The outraged claims of 'funding errors' in the oppositions policies makes me laugh too (laugh being a euphemism for despair), completely over looking, as it does, the 'funded' policy cost of the \$100 million per day added to the budget deficit.

On other thing though, \$43bn with a substantial part spent on paying 'compensation' for the use of something the government already has majority ownership of. Brilliant! Well done the government negotiator on that deal.

What exactly will Telstra do with that \$11bn? My guess is a fair portion will go into 'marketing' give aways, ensuring the playing field remains unlevel for a long, long time to come.

Wait a minute... wait a minute... if, \$11bn buys use of Telstra infrastructure so NBN Co doesn't have to do a parallel build, then, wasn't the majority of the \$43bn ear marked for that, so, what is the remaining \$32bn needed for? That's one hell of a CRM system they are going to end up with.

Incidentally, for just \$2bn more, the coalition could deliver not just 100Mbps but 1Gbps to the NBN footprint households, for \$25 per month. That's my election promise, so vote for me. And hey, it's no more nebulous than anyone else's plan.

Posted by Steve Waddington at 10:07

Friday, July 30, 2010

Early Morning Call

My mobile phone rang this morning at around 5:05am. A call at that time can only mean bad news, and can't be ignored. So various images of network disaster scenarios were going through my sleep-fogged mind as I fumbled for the light switch in my hotel room, then failing to find it, made my way in semi-darkness to the coffee table where I had left the phone, incurring the inevitable, painful, toe stubbing in the unfamiliar room.

As it turned out, it was a NOC agent from a supplier calling to say there was an alarm on one of the circuits - which, I soon worked out from our own monitoring system (once I had powered up my laptop and logged in to our intranet), had been decommissioned several days before. A false alarm and nothing to worry about, thank goodness.

As I was writing to update our own engineers on the call, I was wondering how my mobile number made its way to the primary fault notification list for that supplier anyway. It has been a few years now since I have been included at that level, and I while I knew they had my name and email address from the information provided when I signed the supply contract, I was 100% sure I had never provided my mobile number.

But I must have given it to someone at that company at some time. And from there, it was entered into a contact database, and once a field is populated in a database, who knows where it could end up. Well, cross-pollinated to the technical contact list in this case, for sure.

It reminded me of something similar with Telstra Wholesale six or seven years ago. I had provided a number for SMS notifications for one specific service for a company I had worked at at the time. Somehow that one number migrated its way with the TW databases to all other services, and so I found myself getting 10-20 SMS's a week. Requests to have the number removed worked, partially, for a little while, and the number of SMS's would decrease. But then some system update or something else would happen, and my number would propagate again through the notification database systems and I would be back where I started. It was like some sort of reverse virus.

The thing was, even after I left that company and moved on, I was still getting the notifications, though somewhat decreased in frequency to 1-2 a week, for the next two years.

Eventually, I made an exasperated comment about it to a senior engineering manager in Telstra, and, nice guy that he was, undertook to search the databases manually and remove all instances of my number. Which he did. And it worked.

But then Tom L. was an old school Telco engineer, from days when the PMG, then Telecom Australia was a public utility and far from a corporate predator. Sadly he retired a few years ago (and best wishes to him in his retirement), heralding the end of an era of responsibility taken by a company to its customers it is increasingly rare to see.

So if that is what happens to personal data in, presumably, well controlled corporate databases, what possible hope of integrity is there once your information gets onto a public database like Facebook or Myspace? And I doubt they are staffed by many Tom L's who can take it off if you ask them. Frightening - as people seem to be finding out.

Posted by Steve Waddington at 13:48

Sale and Acquisition - What a horrible mess

I am not referring to the actual acquisition of AAPT, or part of AAPT by iiNet, rather the horrible mess of cables that the iiNet CEO is photographed behind.

Maybe he is trying to look like a 'cool techie' harking back to his sysadmin roots. Maybe he is trying to represent an image of iiNet as a grass roots technical company. Who knows.

What the picture actually shows though is the sort of messy cabling even some second rate backyard startup would find unacceptable. What person with even the slightest pride in their company's networking ability, or the public face their engineering presents, would allow that sort of image to be released to the press?

Planned network builds? Structured cabling? Who needs that? We just bung cables into any old switch we find laying around in a rack.

Ugh. I would be thoroughly ashamed if I was the network engineer behind that mess, and thoroughly embarrassed of I was the Director of a company with a data centre that looked like that.

Posted by Steve Waddington at 09:31

Thursday, July 29. 2010

NBN Map to be revealed

Or did I mean to write reviled? It will certainly be pulled apart and criticized, whatever the case.

But what will it mean?

One way to look at it would be to gauge its accuracy of dates and locations (assuming it gives roll out dates for the coverage areas) by the number of promises so far kept by the incumbent government. Which is zero, which suggests 100% inaccuracy.

A less cynical person might say; based on the likely purchase of Telstra and other carrier assets, in those known areas it is likely to be accurate. So maybe that will be the case. In any event, all will be known, or unknown, tomorrow.

My view is, like all good works of fiction, it will have some elements of reality, shot through with wishful thinking and marginal seat gaining promises. Aragorn will defeat the Dark Lord, ensuring 100Mbps in every school in Gondor, but the Shire will still only get wireless coverage.

If you are going to pick a work of fiction, then it may as well be the best.

Posted by Steve Waddington at 16:11

Wednesday, July 28, 2010

I can see clearly now the clouds have gone

(Johnny Nash)

I was talking to one of our supplier senior managers about, among other things, 'cloud computing' this morning, musing if it would last the way ASP's did at the start of the decade, or if its time has indeed come.

Personally, I am inclined to think the latter. Maybe there is more than a little wishful thinking in that, because the bandwidth and technologies required are appealing to any network engineer. The 'killer app' though is going to be Virtualization (capital V deliberate), that will make servers an on demand commodity.

Think about that for a moment. For normal hardware based servers there is the CAPEX authorization, PO, lead time, staging and commissioning processes that all have to take place. But a virtual server set up can be as simple as picking up the phone, asking for a server instance and jotting down the accessible IP address - all in real time. Or better still, clicking the button on a web page, with the OS option you want, confirming the per month cost, watching your browsers 'wait' icon for a few seconds, and then there you have it, you new virtual server ready to go.

Imagine the flexibility that will give. What I predict we will see is servers being spawned for any new application, database or web site a company needs. And almost certainly many other things I can't even imagine at the moment.

It will also mean a massive leap forward in redundancy available for any sort of server application. When not in use, virtualized server really only take up a bit of disk space - so the cost becomes trivial to create a second, or even third, server instance, mirror the primary server contents, and, should anything happen to the primary server, the backup server can be immediately brought up - or even automatically if that is needed.

No more duplication of full sets of hardware. It all becomes a few dollars a month for some data on a 'cloud' server, and a few dollars more if the CPU cycles are ever needed if a crisis happens.

Kudos to the IT manager who takes that cost saving to the board, I am sure.

Posted by Steve Waddington at 10:58

Friday, July 23. 2010

Suggestions

I may have commented before that a while ago (I am sure that double past tense use in a sentence breaks some grammatical rule) we put in place a 'suggestion box' on the Exetel web site. The comments go directly to John and I, and the polite ones (which most are), almost always actioned in some way, even if it is just to acknowledge the comment was received.

This one today is not untypical:

I was just wondering if you'd consider making the ABC's iView 'unmetered'? I've noticed on iView that it lists some ISPs as providing it as 'unmetered' and I was wondering if you'd consider doing the same. Please, please, please!

and several people have made similar comments in the past.

I find it a bit puzzling though. Because all downloads off peak are 'unmetered', so no one is constrained by a limited number of specifically designated sites to download from. It is also interesting to note that since we introduced off peak plans some years ago, every other ISP has followed suit, and so it seems off peak trumps unmetered in terms of plan attractiveness to users (I am sure a metaphor used twice in a sentence doesn't break any grammatical rule).

Then of course, even on peak, there is no extra cost while within the plan quota.

I guess there is 'free' which is good and so 'even more free(er)', which the advertisement banner of 'unmetered' implies, which has to be better.

More for free! Fantastic! I'm going to get that right after I pay the last installment on the perpetual motion machine that nice man sold me.

I can tell you one thing; unmetered content certainly isn't 'free' to the ISP. Today, (and Exetel almost certainly can't be buying at the best price for global Internet) the past-our-border portion of the total cost to supply an Internet service is in the order of 7-8%. If other people buy better than we do, it would only mean that for them, the ratio is even lower, say 5%.

What that means is the so called 'unmetered' content, apparently 'free' to the end user, costs the ISP around 95% of the full per Mbps of global Internet access anyway.

So, you either believe your perpetual motion machine was a good deal, or, you have to accept that the unmetered content so generously offered by your ISP is paid for somewhere else in the total amount you pay.

Thinking back, it seems to me this is a fundamental error in the cost of service supply that was made by 'someone' years ago (around '98-'99 I suspect), and has been duplicated (or mindlessly followed, however you want to look at it) by others ever since.

In those days, global Internet access accounted for around 1/3 of the operating cost of an ISP, and domestic backhaul added another 15-20%. So by supplying traffic from a local source, such as peering or caching, was an enormous benefit to the ISP's bottom line. Ignoring the cost of peering, the proxy server cost, the engineer/sysadmin time etc, etc is easy to do, as in most cases they are considered necessary anyway, or treated as sunk costs. Therefore at some point, on someones spread sheet to financial management, that content delivery would have shown up as zero. And the marketing guy that saw that would hardly have to be a genius to realize he could be a hero by designing a promotion that offered content that cost zero to the company for 'free' to their customers.

One deception leads to another, and a decade on you have people really believing unmetered content is actually free. No wonder that word is considered the most powerful in marketing. That, and peoples unerring propensity to believe what they want to believe.

Posted by Steve Waddington at 08:32

Thursday, July 22. 2010

HSPA in the Country

From my mate Bill, and rural HSPA pilot site:

May 2009

Towards the end of last year I was talking to Steve about available options to deliver VOIP over a wireless broadband set up and he mentioned that EXETEL were exploring the feasibility of rolling out just such a product in WA with their HSPA package, and would I be prepared to become a test site to trial the system. Saving money on phone calls and staying connected sounded like a good option to me.

To give some background, Sharon and I have owned a small (100 acres) rural property in Napier via Albany in W.A. for some 15 years. Situated 20kms out along the Borden road, about half way to the Porongrup Ranges, our property is split into three where the Yellanup and Napier Creeks meet before departing for Oyster Harbour.

We moved back to Albany from the goldfields in 2005, converted a shed on the property into temporary accommodation, and moved out here in 2007, with a composting toilet but no running water (except the creeks which run all year), no power and no phone. Although power and telephone lines run along our top boundary, the prohibitive cost of running them underground to our preferred building site, and a general desire to explore alternative stand alone power and communication systems led to the decision to go Solar powered and wireless communication, happily the technologies to do this are all happening now.

We have a very basic 12v power system consisting of 4 Trojan Deep cycle 6v batteries, four 64 watt Sola Sun Panels a 20 watt regulator, and a 350 watt inverter (for Lap Top and Printer). We run a 12v Waeco Fridge/Freezer, which is a big improvement on the ice chest we moved out here with, a 12v water pump and numerous 12v low power fluoro lights.

Our lifestyle might be considered on the austere side of luxurious with no TV, no big stereo, and no electric kitchen appliances. There's no washing machine either, which is a pain, but we like the quiet and the space and the birds; it suits us. At night we can listen to frogs and look at the stars or we can surf the net and watch DVD's on the computer, our girls can practice violin without offending anyone (they're actually getting quite good now), run around the farm or do research for their homework and keep in touch with their friends on MSN Live. I walked outside a few minutes ago to get some wood for the fire and the Milky Way is absolutely magnificent, dominating a cloudless sky and not another light around.

May 2010

We are still using the same jury-rigged hardware we started with and it still works fabulously, speed and reliability are both excellent. Not quite streaming video on demand levels all the time, but better than the setups we've looked at in town. A lot depends on the Optus cell being used and this actually improved markedly in the last 6 months or so. For the girls, Facebook has now taken over from MSN live as a means of staying in touch but not much else has changed, which is just how we like it.

Bill Waideman

Posted by Steve Waddington at 08:43

Wednesday, July 21, 2010

Increase in TIO complaints?

Reading this, in particular:

ACMA chairman, Chris Chapman, said the authority wants to find out what is behind the high numbers of complaints being made to the Telecommunications Industry Ombudsman.

My view; caused the by increase in TIO agents, each with a performance (aka revenue) target.

Though it is fairly predictable what the 9 month inquiry will actually give birth to. What a waste.

Posted by Steve Waddington at 11:58

A short Break

From work, and a longer one from blogging. If you don't count traveling to a convention in the US (and having a very nice time there), then I took my first 'holiday', for an extended period, in six years. A few days down south with my sons who were on school holidays, to show them what real camping is about (gale force winds, hail, and below zero nights - ah, the south west in winter), but nothing more beautiful than the sight and smell of dripping, vivid green, pristine Karri forest when the sun breaks through the clouds. Then a couple of days in Albany catching up with friends. A nice refresher to start off the new financial year.

I am very much looking forward to this Financial year. We have expanded our product line significantly in terms of the 'value add' to services. It has been a long road, developing both skill sets and products at the same time, but I feel this year we will be seeing the pay off of that investment.

One of the more significant development is we connected our first Virtual PABX customer at the end of June, and our second will be coming on line next week.

Also, the AI project is kicking along nicely in Sri Lanka. Dr Kumarawadu's small R&D team at the SLIIT has done an amazing job of bringing together the beginnings of a workable framework, and we should be able to begin migrating our knowledge base to it within the next few weeks.

Posted by Steve Waddington at 09:05

Thursday, April 29, 2010

Interop '10, and the forecast is...

... for clouds, clouds and more clouds.

It has been a long, long day today. Starting with jet lag induced 5am sleeplessness - or a good opportunity to catch up on email, then the keynote on Cloud Computing by IBM at 8:30, and a full day at the exhibition listening to what various people had to say.

It is information overload in extreme. It is going to take me a little while to mentally sift through all the information and find the signal in the noise. Starting with the IBM keynote itself.

No keynote tomorrow, but a couple of seminars on Unified Communication I am going to catch. The title of the first one is 'Unified Communications . . . Finally?'. I like the question mark on the end. I have to say though, from what I have seen this year, it is a lot closer then it was last year.

Posted by Steve Waddington at 16:37

Wednesday, April 28, 2010

Interop '10 Revelations

It's amazing the things you learn when you travel abroad. I had a bit of a chuckle at this morning's keynote presentation by Avaya, who have, it seemed, discovered that SIP is a unifying standard for communications - all that proprietary stuff, huh, forget it.

In all seriousness though, I thought the parallels they drew between TCP/IP being the base standard for data communications and SIP becoming the base standard for voice - unified - communications was interesting. And even though there are 27 mainstream flavours of that standard right now, there is no doubting it will be a very good thing when SIP becomes as standardized as IP.

The best one liner of the keynotes was also from Avaya - 'Work is a thing you do, not a place you go'.

The best 'gadget' I have seen so far (not at Interop as it turns out), was yesterday when we called into Fry's for a quick tech purchase top up. One of the things on the list was a pre-paid SIM card with cheap international call rates, because apparently not everyone in our group finds it convenient to use the free Wifi and VoIP at the convention to call home (I know, go figure). I guess I must admit, I could perhaps see the benefit of a phone you could use in case there was no wifi signal.

Instead of a pre-paid SIM though, what we found was a Nokia handset, with pre-paid SIM included, \$20 of call credit at reasonable international rates (using VoIP BTW), in a neat slide off the rack plastic bubble package from My Global Talk.

How much?

Guess. Go on. No, less than that. Ok, I'll tell you. \$25 on special. Nor is the handset locked, and it can be used on any network if you replace the SIM. The handset is pretty basic, with the features common on phones of about five years vintage - though half the weight, half the thickness and twice the battery life. Nothing amazing about the technology, but the price point I thought was something else.

The worst thing about the show so far has been the Xirrus wifi. I am not sure if it's my phone or something else, but the connection, that was rock solid last year, has been dropping in and out all day.

Cisco have used the event to launch their 'Clear Air' wifi technology. Part of the demonstration was that it can identify wifi destroying sources such as web cams, microwave ovens and deliberate wifi blocking devices, isolate and neutralize them. Not that that would have anything to do with the uncharacteristic drop outs from the Xirrus network. Only a paranoid conspiracy theorist would think that. I'm just saying, that's all.

Cloud computing, definitely the hottest topic in US IT land. I am looking forward to the keynote on that tomorrow morning.

Posted by Steve Waddington at 10:11

Tuesday, April 27. 2010

Interop '10

In Las Vegas this week to attend Interop 2010, and looking forward to it. A couple of things different this year; for a start there are four of us making up the Exetel contingent. Also, there seem to be even fewer attendees this year than last year, and last year was well down on the turn out the last time I went before that four years ago.

Maybe that is because, on the early news this morning 'experts' we talking about how there were signs of recovery for the US economy, but other warning of a double dip back into recession. It sort of reminded me of the same sort of financial news were were hearing in Australia eight months ago. I guess an economy the size of the US (which, let's face it is half of the worlds economy) has a bit more inertia that Aus, whichever way it heads.

One pleasant surprise was the Air New Zealand flight we took to get the LA. Last time I flew with that airline, in 2001/2002, the plane was so old it still had ashtrays in the arm rests. This time though, I don't know how old the 747 was, but the cabin fit out was brand spanking new, with state of the art wide screen entertainment, nice food and a very pleasant high end range NZ wine selection.

The exhibition and conference proper starts tomorrow. Cloud computing and visualization are the two biggest topics, and the focus of most of the conference sessions this year.

Posted by Steve Waddington at 02:57

Thursday, April 22, 2010

Beware the Ides of April

It has been, without doubt, the worst week of problems we have had for a long time. While each one, taken as an isolated case, was understandable, taken all together, which is what happened, requires some serious thought, and perhaps even more serious action.

One of the most important things I have learned, which running your own business (or being a parent) brings home to you is; there is only one place 'blame' can be laid when things go wrong. Or, as popularized by US President Truman, 'the buck stops here'. If he did nothing else in his presidency, the least you could say is there is someone who knew how to take responsibility.

The second thing is, never let yourself become complacent. Or perhaps lazy would be another way to say it.

And third, having accepted the above, consider how to fix it and then take action without delay.

One of my favorite quotes of recent times is 'the past is another country and its borders are closed'. No point dwelling on it, but make sure it isn't repeated in future.

So, moving on, I have put a lot of thought into how you can teach that to people who, in their own areas, have to take responsibility. Can you actually teach 'wisdom'? I don't know, but if you can, I think there is a lot to be said for continual repetition.

Posted by Steve Waddington at 15:10

Thursday, April 15, 2010

Hack Attack

A 'hack', to use the colloquial pejorative term is used to describe a writer of low quality, rushed, and perhaps in many cases largely plagiarized work. And we are all familiar with the romanticized 'hacker' of the Internet, lurking being every virtual corner and bush to bring ruin to honest, god fearing Internet enterprise.

I am sure this hasn't been reported particularly accurately or well. Or even written well for that matter - I mean, 'A senior Government source said then'; it doesn't even make sense as a sentence, let alone as a supporting statement from an authoritative source, presumably, in support of the article. A hack reporting a hacker...hmmm...

So just how does a DOS attack on an email filter service one company uses in the US slow internet access to a 'trickle' for that and a bunch of other companies? Very hard to see how that could be. My guess is it's been misreported.

Anyway, the implied scare to national security (a very slow news day no doubt), made me think of a funny incident that happened some time ago in New Zealand.

Coming back from a trip to the US, my passport stayed in the back pocket of my jeans and somehow ended up on the wash when I got home. Hardy documents passports, it survived pretty well, but obviously needed replacing. Which I planned to get around to one day.

A few weeks later, for some reason that was no doubt very important at the time, I had to go to New Zealand the next day to meet with some people interested in investing in the company I was working for (or something like that, as far as I can remember).

The photo and ID page on my battered, but clean, passport was in tact. There was a little bit of warping as the paper dried out, and the edges were a bit tattered, other than that, as a document, it was still completely legible. No choice any way, with less than 24 hour notice, no change of a replacement, so off I went.

Passing through customs in Australia, the official took one look at the passport and said 'Been through the wash mate?'. I smiled meekly and nodded and he waved me through.

Different matter in New Zealand though.

'What's this?' said the kiwi customs officer (actually she said 'Whut's Thus', which I took to mean the same thing).

'Went through the washing machine' I said, expecting the same sort of nod of understanding and pass through as he Australian counterpart had given.

But no such thing would happen in New Zealand. No sir-ee. This is post 9/11 you know, and you can never be to careful about these things. Especially in somewhere as important as New Zealand. I could see the suspicion growing on her face with each wrinkled page of my passport she turned.

'Stund uver hure pluz Sur' I was instructed, moving toward the square of yellow tape on the floor she indicated, hoping I had understood what was clearly an instruction to do something, which I guessed was to stand there.

Once I was safely within the terrorist escape proof yellow taped area, the officer went off though the 'Security Staff Only' door, and for the next twenty minutes I saw the occasional glimpse of different people giving me various degrees of scrutiny through the glass pane in the door.

Eventually, after who knows what deliberation must have taken place, she returned, and told me 'We can't accept this as a valid passport. You will have to get a replacement'. What a bummer I thought, they wont let me into the country, a completely wasted trip.

So I asked, naturally enough, what I was to do now. Do I go back to Australia?

'Oh no' I was then told 'You can't do that because as far as we are concerned, without a valid passport you don't have a

country of origin, and therefore you can't go 'back' anywhere'.

Which, I have to say, took a bit of digesting, and I had always thought of New Zealand, despite the atrocious accent, poor rugby and cricket skills, and chip on the shoulder, as not completely antagonistic to Australians. But now, well, where exactly does a person of with no country of origin and no passport go?

I was about to find out as she continued, saying 'We can let you in, but you will have to get your passport renewed at the Australian consulate to be able to leave'. She didn't stamp the passport (which wasn't valid anyway), but just waved me through the gate and off I went.

So there you go. Draw your own conclusion about how desperate a country has to be to improve their sheep/human ratio to have that sort of swinging door policy for entry.

Which isn't the funny part of the story. So I continue...

The next day, we had whatever meeting was so important, and my two traveling companions (an Irishman and an Englishman as it was) and I went off to find the Australian Consulate, a couple of blocks away from our meeting in downtown Auckland. I went in, explained what had happened and the helpful staff person there took my nominal fee, the name of an Australian Citizen referee they could use to verify who I was, and told me to come back after 11am the next day to pick up a new passport. No problem at all, almost too easy.

Next day, same drill, meeting, then the three of us dropped by the consulate (since it was on the way to lunch anyway), as I picked up my passport. I only mention this because the fact there were three of us makes no difference, other than 'an Irishman, and Englishman and an Australian went into a...' is a funnier way to tell a story.

'Your passport isn't ready yet. If you want to take a seat it should only be ten minutes'. I was told.

Just at that moment, I got a call from the Perth office telling me of some serious problem that was happening on the network. This was in the days when mobile phones were just phones, and the nearest computer I could use, short of trying to find an Internet cafe, was half an hour away back at the hotel.

Or, there were a couple of computers in the consulate foyer where we were waiting, with a web page showing immigration information to Australia. A quick look at the terminals showed they were locked down so the only executable on the window was IE. Clicking the 'Start' button showed the same thing - except - the run field bar was also active.

I typed in 'command', and what do you know, DOS prompt. 'Telnet x.x.x.x' gave me a Cisco login prompt. Four telnet windows and a few minutes later I had found the problem was a denial of service attack (a real one, not an excuse for a network operator error), and I had put the filters in place to kill it off.

Just as the consulate person came back and told me my passport was ready, handing it to me across the counter.

My two companions thought it was all a huge joke, and were doing their best to look as suspiciously unsuspecting as they could. For my part, I casually shut down the telnet sessions and closed the DOS windows giving a polite nod towards the security camera as I finished.

With the glass door just closing behind us as we left, we saw the 'staff only' door to the back office open and a puzzled/worried staff member strode into the foyer area to, perhaps, we thought, determine what threat to national security was occurring from the consulate computers. I mean - who knows what sort of people could be wandering around the streets of Auckland, let in with soggy passports cut out from the back of a wheatbix box.

Posted by Steve Waddington at 16:20

Wednesday, April 7, 2010

The Caches Three

Current bandwidth generation for our three PeerApp caches:

Sydney

Melbourne

Brisbane

The traffic offered to the caches is 3Gbps, 800Mbps and 500Mbps respectively.

Posted by Steve Waddington at 02:36

Thursday, April 1, 2010

Melbourne Cache

The upgrade work at the Melbourne POP was successfully completed yesterday, which included a new Cisco 10000 LNS router and the installation of a PeerApp cache.

The cache was activated last night and we are currently seeing a peak of 30Mbps, average 20Mbps generated traffic. As expected, it will not deliver its full benefit until the hard disks fill up, which takes about a week.

Victoria, for some reason, continually consumes 50% more bandwidth per user than any other state. We have never figured out just why that is. So it will be interesting to see how the cache generated traffic compares to that of Sydney and Brisbane.

Posted by Steve Waddington at 11:11

Network Shutdown

Well, it had to happen at some time, and now is the time.

For the last month we have been testing IPv6 with staff accounts, and out one bandwidth supplier that can support it. It all went pretty well, most modems we tested had automatic support. Some required a manual change to activate it, and less than 30% didn't work at all.

So today is the day we make the change.

At noon today the network will be shut down for 30 minutes as the IPv6 code change is made. When it re-activated, the whole network will be IPv6 only.

Posted by Steve Waddington at 10:29

Wednesday, March 31, 2010

The Tyranny of Distance...

... is to be replaced by the tyranny of Labour's communications policies:

There is no worse tyranny than to force a man to pay for what he does not want merely because you think it would be good for him.

and

The whole principle is wrong. It's like demanding that grown men live on skim milk because the baby can't have steak.

(both quotes attributed to R. A. Heinlein)

What could better sum up the policies of a \$43bn NBN and a mandatory web filter?

I don't know if many people have looked at 'chatroulette'. If you haven't, don't bother, I certainly regret looking at it to see what the reported 'new phenomena' in social networking was about. It is pretty much rock bottom by any measure of any aspect of humanity. I wonder how the mandatory filter will treat that site?

Unlikely either policy will change though, as another Heinlein quote points out:

Never underestimate the power of human stupidity.

My predication is that, as the nations wealth is squandered on handouts, destroying the health system, censorship simply for the sake of censorship and white elephants no one needs or wants, the wreck that is left of the economy will be blamed on the bad luck of 'external factors' 'beyond our control'.

Throughout history, poverty is the normal condition of man. Advances which permit this norm to be exceeded " here and there, now and then " are the work of an extremely small minority, frequently despised, often condemned, and almost always opposed by all right-thinking people. Whenever this tiny minority is kept from creating, or (as sometimes happens) is driven out of a society, the people then slip back into abject poverty. This is known as "bad luck."

Heinlein again of course.

Posted by Steve Waddington at 16:09

Friday, March 26. 2010

Brisbane CACHE GRAPHS

These are the graphs for the Brisbane PeerApp cache.

`Weekly' Graph (30 Minute Average)

`Monthly' Graph (2 Hour Average)

So far, we are hitting about 60% of what we expected for the volume of traffic offered to the cache. Our experience though is that this will continue to improve over the next month to where we see a peak generation around 150Mbps and average of 70Mbps.

ps. The Windows Live Writer blog editor doesn't seem to be able to handle uploading files very well.

Posted by Steve Waddington at 10:57

Tuesday, March 23. 2010

Perth Storm

Not the Netball club, but the torrential downpour, thunder and gale force winds we had in Perth last night. The light show we had from our patio looking north to the city and east to the hills was spectacular. Predictably, with power flickering on and off during the height of the storm, our ADSL connection was still down for about an hour and a half after it passed. Though I imagine other people, 90,000 still without power the next morning apparently, would have experienced much longer outages. And I wonder how many had some electrical damage to their computer/modem as a result? An interesting stat is that faults reported from WA were less for a Monday/Tuesday than in previous weeks. I would have thought it would go the other way, but I guess most people are sensible enough to know that an ADSL problem at during the worst storm in 20? years is almost certainly going to be the cause, and that these sort of things get fixed in time. Anyway, like anyone who knows that copper is a pretty good conductor of electricity, all of our household technology was unplugged for the 2 hour or so duration. Two hours without Internet access! However did you manage? You ask. A laptop and HSPA of course. What a lifesaver. Logging into the Perth POP, I could see nothing untoward there, and the decrease in connections and data use in WA was local to individual users/exchanges. Connections gradually dropped by around 25% between 6pm to a low point at 9pm, then slowly picked up until they were back to normal (for the time of day) by 11am this morning.

Posted by Steve Waddington at 16:54

Friday, March 19, 2010

History Lesson

I see Telstra are taking a bold step into the '70's.

Put 'HD' in front of any technology, no matter how old, and it has to be a sure winner, right?

Such innovation!

Meanwhile, we have commenced development of our own integrated voice/data/video solution, to include video conferencing. Which will be based on open source applications and, a key criteria, will integrate with an Asterisk PABX.

I don't know that it will be HD, though that will be a function of the camera. What we aim to end up with is a product that will be practical for an SMB size company, and easily affordable as an add-on or part of a bundled solution.

Posted by Steve Waddington at 12:59

Thursday, March 18, 2010

NBNCo - Is there anybody in there?

Email sent today to info@nbnco.com.au:

Dear Sir or Madam,

I have tried to contact you via your web site on Monday and by phone yesterday. After taking my details, your receptionist told me someone would 'definitely call you back soon'. So far there has been no response to either.

I am trying to contact you to determine if and how Exetel can become a reseller of NBN products and services. I would appreciate the courtesy of a reply.

Regards,

Steve Waddington
Director
Exetel Pty Ltd

Maybe there is no CRM system in place yet. I am sure one can be found for a few hundred mill or so that will do the job.

Posted by Steve Waddington at 11:56

Monday, March 15. 2010

Holy Blow-out Batman

How many NBN administrators does it take to change a light-bulb? Well, according to the latest news about the NBN cost, that question can't be answered until HUNDREDS OF MILLIONS OF DOLLARS have been spent on an 'enterprise management system'. That's not the staff, cable or anything else to actually do with running the network, that's just the system to, essentially, run the payroll.

Bugger technology, stuff any advance in productivity, screw modern best management practices. What is abundantly clear that is being built here is a good, old fashion, pre Keynesian bureaucracy - of staggering proportions.

Seriously. You put fiber in the ground. You connect it to some hardware. It works. If it breaks, the contractor that has the supply and maintenance contract fixes it - NBNCos doesn't have to do anything. So too with the hardware, a Cisco, Nortel, Huawei or whoever will look after all that under the terms of supply.

A good, modern network should be a very inexpensive thing to run. We are not the PMG in 1910 any more. They even make the point they are 'starting from scratch', and don't have legacy systems, or retail customers, to consider.

So has anyone actually done a rational cost justification of this? Apart from Accenture that is - to whom I must say - well done fellas, someone is going to make partner from that deal for sure. No bonanza???? A bloody big one for Accenture if you ask me.

And the quoted justification these genius's come up with for the cost? It's not as much as the 'billions often spent by tier-one retail telcos'. It's like comparing the a 400kg person to someone who weighs a mere 320kg and saying 'It's ok, I'm not as fat as that guy'.

Another question I haven't seen anyone ask, or answer, is; just how much is this going to cost to run? My assumption would be 'not much' - perhaps in the order of a few percent per year of the build cost. But when you see information like this, that assumption really has to be questioned. Are we also going to be saddled with a hugely expensive ongoing maintenance cost? Just how much is that going to be?

Quite a lot if the price of the management software to handle a hundred or so customers and a handful of contractors is a 'few' hundred million dollars.

Tell you what, since Exetel may well be one of the customers, just give me my three or four million dollar share of that now (or more likely 30-40 million), and I'll forgo the ten years of heartache and hassles that will inevitably result.

Posted by Steve Waddington at 09:42

Sunday, March 14, 2010

Brisbane Cache

The PeerApp cache engine for the Queensland POP went live on Friday morning. Installation required a little bit of rack re-arrangement, and about 5 minutes downtime to change circuit route paths. So overall it was a well planned installation and successful deployment by the network team.

By Friday evening, the cache was only generating just over 1Mbps of traffic, however as I check just now, I can see it has peaked at 40Mbps and averaging 25Mbps.

With 700Mbps of offered bandwidth in Brisbane, we should be able to get to at least 100Mbps average, and maybe 200Mbps peak, I should think. It will take a week or so for the hard drives to fully populate, and possibly a little tweaking to get the most out of it, if it follows our Sydney experience.

I'll post the MRTG graph once the link is up for it next week.

Posted by Steve Waddington at 19:02

Friday, March 12. 2010

User Experience effect of Peering and Caching

On Monday, with little fanfare we completed a setup for direct peering with Google in Australia. Incrementally to the overall network, it hasn't made a lot of difference, but has, never the less, improved traffic flow through that part of our network a little. Which is always a good thing, and all these little things add up.

So last night while working back in the office, I was updating my RSS feeds to some of my favorite podcasts. The office LAN, since it connects to our data center directly, effectively has a 400Mbps connection to our Internet backbone. Though for all intents and purposes, since my computer NIC connects at 100Mbps, and the computer IO itself can only sustain 40Mbps, I have an effective connection of 40Mbps.

Anyway, the point is, I downloaded a couple of the Science podcasts from the ABC Radio National site. Not very big files, in the order of 10-20Mbytes. The usual Firefox download window popped up - for about half a second - and then said 'All downloads complete'.

'That has to be an error' I thought, it must have downloaded a zero length file or a null link. So I did it again - same thing. Then I looked for a bigger file and found the Jan 16 Science Show. Again, the download window blinked for maybe half a second before reporting the download had completed. So I went and looked for a video podcast, and at around 90Mbytes, the Gardening Australia show took all of 3 seconds to download.

I don't know why I am so amazed by that, I essentially had a wire speed connection to the ABC service through GE peering 'just down the road' at our Sydney datacentre. Perhaps it was because the Internet download speed was noticeably faster than the data transfer from my laptop to my USB 2.0 mp3 player. Nice to see that all that 'extra' peering and caching we have invested in pays off in such a direct way though.

Posted by Steve Waddington at 06:47

Wednesday, March 10, 2010

Test Post for Windows Live Writer

My year old HP laptop died yesterday after a couple of days of blue screening, black screening, complaining it couldn't find its HDD and memory errors. Its replacement is a Dell running Windows 7. I have to say, I was not sorry to see the back of Vista, and, while not looking forward to having to make the change, since it was forced on me anyway, interested to see how Windows 7 would live up to its advertising hype. Functionally, 24 hours in, all I can say is that the annoyances of Windows 7 are slightly less than Vista, and a little easier to turn off and reconfigure. I will have to see how it goes over a week though, to be fair. One thing I found while disabling whatever that panel thing in the middle of the screen was, is this Live Writer utility. It doesn't have native support for Serendipity, but a quick google search for "serendipity windows blog writer" found this useful post on the Serendipity site. A minute or so later, and here I am, using it now. It seems to work fine, and the editor is a little better than the default one in Serendipity itself. Now to click "Publish", and if you are reading this, it worked.

Posted by Steve Waddington at 17:55

Tuesday, March 9, 2010

Professional Services

We had the first meeting yesterday to discuss how we could provide a database re-write for a medium sized businesses inventory system. Not a very startling or revolutionary act in itself, but it is significant in that it is the first time we have looked, formally, at providing software development services outside of our own need.

Raymond presented his draft proposal, estimating 2,000 hours of development time would be needed to completely re-write a system that had been in use for the last 15 years. The current system the end user has was written in something, if I had ever heard of it, I had forgotten, and the task at hand was to duplicate the functionality using MySQL. All in all, about a six month project.

The first thing that stuck me was that the proposed cost seemed very low. And by very low, I mean ridiculously low. I don't know what the 'going rate' for software/database development is nowadays, but I remember from 10 years ago when I was at CGE&Y, the lowest rate on the card, for the most junior developer, was \$120 per hour. Two junior programmers need one more senior supervisor, and the project would be overseen by a manager or senior manager, so the proposal would be something like:

2 x developers @ \$120/hr for 2,000 hours
1 x Snr developer @ \$180/hr for 2,000 hours
1 x Manager @ \$240/hr for 400 hours

About a million dollars in other words.

I also recall at the time, some consulting companies had started hiring in India, which, I heard, would allow a reduction in developer billable time to \$90 per hour, or an effective price-to-the-customer reduction of 30% for any particular project where those resources were used.

So I looked at Raymond's figures, reworked them, thought about it some more, ran them again and, well, still found it hard to pass the 'sanity check' of what I thought such a project should cost. No matter how I tried I just couldn't get anywhere near the \$500,000 mark, let alone \$1M - well, not without just multiplying the final figure by some arbitrary integer value.

I guess I will just have to face it, we really do have a very, very low cost base for development. Which does make sense in a way, because it is the only way we could have afforded to develop the systems we need that have allowed us to survive, and complete, in this tough market place.

Then why not make those efficiencies available on a commercial basis too?

The price I was struggling with is, by the way, \$150,000. Crazy, right?

Posted by Steve Waddington at 09:34

Monday, March 1. 2010

Good Grief (Charlie Brown)

I don't know what Charles Schulz had in mind when he coined that phrase, but it always seemed to me like an expression of disbelief at the eternal optimism shown in the face of crushing failure that was the inevitable result of whatever Charlie Brown did.

What I am finding is it is very hard to write anything positive at the moment, as the stupidity in relation to anything to do with Australian government expenditure plumbs ever new depths, from the likes of:

- * Conroy - \$100M here, \$100M there. No Tasmanian will be without a service they don't want and can't afford.
- * Rudd - with the tenacity of a cockroach to never let being completely wrong deter you from a course of action.
- * The Minister for Electrocution and the Destruction of the Environment - "botched scheme was developed during difficult economic times." Ah, well, completely understandable then.
- * Wong - Was any thought put into the previous decision? Has any more thought been put into this one? The sum total seems to be 'well that didn't make us as popular as we thought, so let's do this instead'.
- * Gillard and Prof. McGaw - Hope? Hope?? Shouldn't there be a little more certainty than just wing and a prayer hope that a revised curriculum will work?

Where is the outrage of the 'free press' at all those contradictory, wasteful and pork barreling policies? Or even just the basic job of questioning them?

I don't know, I just feel this growing numbness of the inevitability that Rudd's determination to squander the nation's wealth will succeed. All the 'promises' seem just like Lucy's to never pull the ball away again before Charlie kicks it.

And tomorrow will be a new day, with a new sun, no doubt.

Posted by Steve Waddington at 14:01

Friday, February 19, 2010

Managers Conference Results

One of the benefits of giving a group of people a lot to do in a very short time is that it highlights strengths and weaknesses of both the individual and the group. In a normal working week it is rare, if not impossible, for 12 busy people to be able to sit in the same room for an extended period of time and nut out the things that hinder or annoy them.

So when an environment is created that causes those two things to happen it is impossible for people not to learn something about themselves and others, and, with a little attention to the goals of the course, ensure a number of positive results are the outcome.

This was by far the biggest budgeted course I have run, held at a Star City function room with hotel accommodation for all 12 attendees. We kicked off on Friday evening with a few drinks and dinner at Astral, then (I imagine) everyone went back to their rooms to study the course notes for an 8:30 start on Saturday morning.

I have to say that the facilities provided by Star City were superb, and the catering went even beyond that, to whatever place descriptions go when praise becomes hyperbole through lack of skill on the part of the writer. Fantastic job Star City Functions - if there is an industry award for excellence, you deserve to win it.

I find 12 to be about the perfect size for an event like this. For team exercises there can be three groups of four, for individual exercises everyone can present their work within a reasonable time, and there are enough people to give a wide range of views in discussions.

Within the main theme to 'improve written communication', there were five main goals to:

- * improve written communication skills and comprehension of Exetel managers
- * impart knowledge of department operation and requirements of other departments between managers
- * reinforce Exetel as a sales company requiring immediate actioning of tasks
- * engender the concept of decision making and risk
- * impart an understanding of the financial responsibility for a management position

and four deliverables:

- * A standardised, concise and clear reporting format
- * A standardised capex and opex proposal format
- * A published Zero defect process for each department
- * Accurate capex and opex forecasts on an on-going basis

I also had a mental list of about 20 minor goals to be achieved, that I would steer towards as the opportunities arose (as they always do) during the course.

I don't know how other people do it, but for me, running a course is always a matter of a continual feedback loop. You never know how people are going to respond to information (though after a while you can get a pretty good feel for the way it will go), so the presenter/coordinator must always be alert to the way people are reacting, and adjust the information to have the most effective impact, seize an opportunity to drive home some point, or even be prepared to change the agenda if that will better serve the purpose of achieving the objectives.

Seventeen presentations were given by myself, Tim North (from Scribe Consulting) and senior Exetel managers during the course. Each one was rated against four criteria on a scale of 1-5 by each attendee. With '1' being best and '5' being worst, the average rating for all presentation was 1.5, and no presentation scored lower than 2.5 on average.

There were also four group and three individual exercises. The general flow being; get some information - apply the information to a task - present it for peer review - revise it if needed - arrive at a result. Which I find an effective way to do in a few hours (where uniquely everyone is in the same room that can contribute to the outcome) that may take days or weeks, or even never, achieved otherwise.

Two and a half days of intense observation, predictive analysis and micro manipulation left me feeling pretty well drained by wrap up time on Sunday afternoon (and only a five hour flight before I was home). Never the less, I was pleased with the results, having achieved the goals, produced the deliverables, and learned some things, good and bad, that will lead to future improvements at Exetel.

Posted by Steve Waddington at 10:34

Thursday, February 11. 2010

When will they ever learn?

(P. Seeger)

In relation to AI, I read this article today, where a number of experts predict when an AI will pass the Turing test, equal, or exceed human intelligence. My view is I tend to side with the opinions to the far right of the graph, though for reasons I am sure not many people would agree with (arising from the nature of intelligence and the complexity of sensor input processing required to generate it, where I take a view similar to Douglas Hofstadter - I think).

Anyway, our own AI program has no such grand ambition, nor has it made any progress at all for quite some time now. And the reason for that is due to our own 'complexity' stumbling block - of no efficient way to edit and arrange the knowledge base, coupled with no spare development resource to develop that. It also became apparent some time ago that such a project, though the promise if successful is high, could easily consume all our development resource, and was probably going to be beyond the scope of what we could 'risk' to commit to it.

That being the case, the project scope never the less would fit very well into an academic research program, where there would be a) keen interest, b) 'many hands' in the form of students willing to undertake programming tasks as part of their course/graduation work c) ability to manage research required than a company like Exetel could maintain over the longer term and d) development of a commercially viable product.

Twelve months ago (February 16 2009 to be exact), I contacted the SLIIT and began discussions to engage in three programs we saw as mutually beneficial. These were:

1. A scholarship program
2. An internship program
3. Endowment of a research chair in AI

We initiated the scholarship program in mid 2009, fully funding the fees of the first two students, which will grow to eight students by 2012. The first two students had excelled in their first year, and we look forward to adding two more to the program this year.

The intern program has also been successful. We have given paid internships to four final year students who have completed the assigned work, and we have been able to offer full time positions to all four. That program also will continue into its second year.

The success of the first two programs, and the development of a good working relationship with the SLIIT over the last 12 months gave us the confidence to proceed with the third. Which is the greatest commitment in both in terms of money and logistics.

SLIIT advertised for the position of Endowed Professor late last year. Though like any very senior position, such a thing is not a trivial exercise to fulfill, and no candidate was found.

I discussed the outcome with Professors Karunaratne, Gamage and Kapurubandara (President, CEO and Dean of Student Affairs) while I was in Sri Lanka last month, who suggest that we expand the range of possible candidates to include a Senior Fellowship option as well. An eminently sensible idea.

This is the current ad being run:

My feeling is Sir Arthur undersold his adopted country making Dr. Chandrasegarampillai from India the creator of HAL. We will have to wait to see the outcome of that.

Posted by Steve Waddington at 11:10

Wednesday, February 10, 2010

Nobody's Fault but Mine

(Page/Plant/W. Johnson)

I am completing the final touches to a series of presentations for a managers training and planning course I am running this weekend. With 12 attendees, each responsible for some department or business unit, it brings home as much as anything how the company has developed in the last year.

The focus of the course for the 'training' objectives is to improve written communication skills. But, as with any event like this, or a good Seinfeld episode, there are multiple layers to the goals and outcomes. A number of diverse threads of 'input' information will be given that broadly follow the main theme, while introducing other concepts and considerations that each individual can pursue, to fit together the the concepts in a way that makes most sense to them.

The goals for the course attendees are:

- * improve written communication skills and comprehension of Exetel managers
- * impart knowledge of department operation and requirements of other departments between managers
- * reinforce Exetel as a sales company requiring immediate actioning of tasks
- * engender the concept of decision making and risk
- * impart an understanding of the financial responsibility for a management position

Which are in addition to what I am looking to get out of it. Every time I run a training course, something that always surprises me is just how people put the pieces together to arrive at unique, and very often far better, solutions than I could have ever thought of.

I see that about 100 people a day look at this blog, so at least that many people know I am far from the most competent person to talk about improving writing skills. So for the first time I am bringing in an outside presenter, Tim North of Scribe Consulting, who will be conducting three presentations on:

- Business Writing for Reports
- Business Writing for Proposals
- Business Writing for Results

I really have to stop using the word 'training' though. It implies imparting a skill or knowledge that someone doesn't have. And someone that doesn't have the skills and knowledge to achieve all of the goals above would/should never be a manager in the first place. The only real 'training' anyone can do is what you teach yourself as a result of first hand experience.

Guidance, mentoring, knowledge transfer, are all more accurate words, though sound either clumsier (or wankier) to use to me.

One of those many threads of information I talked about will be to reiterate that, as a manger, there is only one place to look for responsibility for success or failure of any project.

Never the less, as a result of:

- defined deliverables each attendee to responsible to achieve
- explanation of best practice techniques
- a mosaic of information that may be helpful
- interaction through set tasks

I am looking forward to a positive outcome for each person attending, and a long term benefit to Exetel.

Posted by Steve Waddington at 10:24

Tuesday, February 9, 2010

NetSense Automater...

... is possibly the single most useful tool for any Cisco based network. And we have made extensive use of NetSense scripts for the last five years, using them to automate just about every repetitive task on our network.

This has included some complex operations such as load balancing of circuits, the application of time of day access lists, as well as mundane (but of life saving importance) tasks such as twice daily backups of running code on all our routers. The scripts themselves are straightforward shell scripts run from cron jobs that call the various NetSense utilities. Here is an example from a set we used a few years ago:

```
# run the bandwidth control auto update script

25 16 * * * $HOME/net-sense/templates/control-scripts/auto-do-yellow-orange-remove-tasks.sh #1>/dev/null 2>/dev/null

# make the fairplay pool and load it

20 23 * * * $HOME/net-sense/templates/control-scripts/make-fairplay-cmd.sh #1>/dev/null 2>/dev/null

# add yellow, orange, red group traffic controls

03 08 * * * $HOME/net-sense/templates/control-scripts/add-rate-limit-tasks-yellow.sh #1>/dev/null 2>/dev/null

02 00 * * * $HOME/net-sense/templates/control-scripts/add-rate-limit-tasks-orange.sh #1>/dev/null 2>/dev/null

01 00 * * * $HOME/net-sense/templates/control-scripts/add-rate-limit-tasks-red.sh #1>/dev/null 2>/dev/null
```

And for backups, still in use:

```
14 19 * * * $HOME/scripts/bin/get-router-config.sh
```

The shell scripts (no reason to use bash, other than my limited programming ability) looked like this:

```
#!/bin/bash
#
#
cd ~/net-sense/templates
~/net-sense/bin/copy_to_tftp -rf all-routers -ipaddr 220.233.0.25 -pw logins.var -tftpboot ~/tftpboot -subdir `date
+%y%m%d`-1
```

where copy_to_tftp is the NetSense utility, all-routers is a text file with the ip address of each router, and logins.var matches the router IP to the login passwords. The router config is then copied down to a file tagged with the date stamp of the backup. So simple, and so useful.

The time of day traffic controls we used to use for different end user groups is even simpler:

```
#!/bin/bash
#
#
# add rate limit to CORE router interfaces
#
cd $HOME/net-sense/templates
$HOME/net-sense/bin/config_devices -log ~/logs/rate-limits_update.log -rf CORE-ROUTERS -pw logins.var -cf
add-core-rate-limit-orange.cmds
```

Another script from the database would take the opt-in or excess usage information from user accounts and generate a list of IP addresses that the server would fetch each day, which in turn were used to create the .cmds file used by NetSense. The cmds file contain just straightforward IOS commands that the config_devides utility writes to each router in the CORE-ROUTERS list, and looks like this:

```
interface GigabitEthernet0/2.30
rate-limit output 1000000 10000 20000 conform-action transmit exceed-action drop
```

The rate to apply being based on the number of IP address or whatever other parameters were needed, and then generated automatically by the script.

While another takes the IP addresses and parses them into an access list so the content of 'add-core-rate-limit-orange.cmds' looks like this:

```
no ip access-list standard orange
ip access-list standard orange
permit host 220.233.4.x
permit host 220.233.11.x
permit host 220.233.25.x
permit host 220.233.9.x
permit host 220.233.14.x
permit host 220.233.13.x
permit host 220.233.4.x
permit host 220.233.9.x
permit host 220.233.19.x
.
.
.
!
deny any
```

The access list itself is not used within the traffic shaping. Rather, it is used to determine which path the traffic takes within the core network. Which in turn is achieved by VLAN circuits and IGP announcements - the access list is used by the IGP (EIGRP in this case) which is given a metric lower than the default metric of the network.

This not only automated a time consuming task, but also allowed us to run very sophisticated time of day bandwidth controls on 7301 and 7206 class routers at full line speed traffic volumes (approaching 1gbps), while keeping the router CPU to sub 60%.

Posted by Steve Waddington at 11:07

Monday, February 8, 2010

Cache v2.0 PeerApp progress

The upgraded cache was installed just before Christmas, and production traffic progressively added over New Year and January.

Our biggest problem so far has been to reduce the amount of asymmetric traffic traversing the cache switches.

In a nutshell, the PeerApp system takes IP traffic from all our supplier circuits, decides, via port number, which traffic is cache-able, and redirects that traffic to the cache servers. The cache servers then perform further analysis of the traffic, storing what can be stored onto disk, or delivering what can be delivered off of disk, and passing through the rest of the traffic back to the network.

Very simple in terms of any 're-engineering' we have to do to accommodate the cache - which is to say, virtually none - just plug the external circuits into the cache switch rather than the border router, and then take the same number of circuits out of the cache switch and connect them to the interfaces the external circuits occupied previously.

However, our traffic is not as simple as that. Because, as we have built the network out, we have done so with a blend of centrally supplied (from Sydney) and locally supplied (in each state) bandwidth, based, at each point in time, on the most cost effective delivery method.

The net effect was approximately 20% of traffic through the Sydney border router was asymmetric. That is; inbound packets arrived on a different router to the outbound packets.

While this causes no practical problem for user Internet access, for the PeerApp cache it reduced the effectiveness of, particularly, port 80 caching - which is the main area of improvement we looked for with the new cache system.

Heaven forbid the life of a network engineer should ever become boring. And over the last four weeks our core network team has progressively removed all but a tiny remaining fraction of asymmetric traffic from the NSW network.

End result?

Traffic spike coincide with user plan peak periods.

The new cache is now generating twice the traffic of the system it replaced. With 2.2Gbps of traffic offered, we see peak generation of 850Mbps, and an average of just under 400Mbps generated traffic.

The graph peak coincides with the start of our 'off peak' or uncharged download period. The fast decay of the traffic is, I believe, due to the quick delivery of cached content that is requested. Which in turn matches the 'user experience' of near wire speed (6Mbps in my case) of data delivered directly from the cache.

The New Cache System generates twice the traffic as traffic symmetry has been improved.

Posted by Steve Waddington at 11:39

Friday, February 5, 2010

A win for inet in court...

... will lead to an ultimate loss for everyone else, especially the ISP industry.

Why do I think that? Because what Justice Cowdroy's judgment says to me, and from the look of this, the rest of the world as well, is that the Australian ISP industry has proved itself unable to take reasonable steps to behave like a responsible corporate citizen, and, if the industry won't be responsible for itself, then a change to the legislation is needed to force it to do so.

I thought Cowdroy's admonishment of the inet testimony during the trial was a clear indication, and warning, that 'playing dumb' was not a good idea, and that some future competent prosecution could have a field day with.

So here it is; on the Lateline story last night Neil Gane said "We are confident that government will not intend a policy where rampant copyright infringement is allowed to continue unaddressed and unabated across the internet.". Which to me indicates the court case was only the first battle for AFACT in the war of copyright infringement. Having proved ISP's won't take any responsibility themselves, and the current law (subject to whatever appeals may or may not happen) can't force them to, they will lobby for the legislation to be changed.

Does no one play chess any more or understand what a gambit is?

I guess whatever will happen will happen. A pity though, it would have been so easy with a little less ego and taking a little more responsibility to have arrived at a far better outcome for everyone.

Posted by Steve Waddington at 11:13

Thursday, February 4, 2010

The Best Mobile Broadband Coverage in the World..

.. could it be Sri Lanka?

No kidding. Deep in the jungle in the Bay of Bengal, I was checking email, accessing the company intranet and had a measured download speed of 1.6Mbps.

I can't say Sri Lanka is the best for mobile broadband coverage, since I have never been to Madagascar, Myanmar or the Congo, but I have been to Pemberton, Wagin and Gnowangerup, and the coverage in Sri Lanka is better than there.

The Heritaance Hotel at Kandalama, 4 hours drive from Colombo, 2 hours from the nearest A road, and half an hour on dirt tracks to get to. Mobile broadband speed was 1.6Mbps and rock solid.

Apart from coverage what is the difference between 3G in a 'third world' country and Australia? Well, in Sri Lanka, the much lower per capita income means the carriers can only charge \$5 per month (500 rupees, which is close enough to A\$5 at the current exchange rate) for the service.

A very interesting number, I thought, because if you think about it, it has to be an honest reflection of the true lowest cost such a service can be offered to the retail market, without wanker bundling and/or gouging download pricing.

Communications towers near Pinnawala. Beats the coverage at Pingelly by a country mile.

Ah, but, you say, Sri Lanka has much higher population density and therefore a much better economy of scale! A much touted argument, so much so that like any cliché who ever questions it?

So let me tell you, six hours drive to get to a place is still six hours drive - whether the distance traveled is 600km or 200km - the effective cost is still the same. Sure Sri Lanka has a much higher population density, but higher per subscriber per square km than, say, New South Wales or Victoria? It would be hard to see how.

Less competition? I don't think so. There are three competing networks in Sri Lanka too. Lower cost of bandwidth? Doubt it. I don't know what a good wholesale price for bandwidth is in Sri Lanka, and maybe we don't buy very well there. But I do know that the price we pay per Mbps, taking the best quote from a number of suppliers, is around the same as the price in Australia in 1999.

Maybe its just me, but if a government was determined to spend \$43bn on national broadband infrastructure, wouldn't it be preferable to have a goal of making access cheaper and more widespread, rather than more expensive and only for cities and large country towns?

Dawn breaks over the jungle near Dambulla. A good time to check email - because you can.

Mr Walker, perhaps you need to relocate your skull cave to where the 21st Century pirates are.

Posted by Steve Waddington at 10:30

Thursday, January 28, 2010

The Elephant and the Farmer

Last week I was in Sri Lanka, part of the trip agenda to see the site of the Model Dairy Farm Project. The site, in the so called 'dry area' Matale district (still looked like lush rain forest to my WA wheatbelt adjusted eyes), is six hours drive - 200km - from Colombo.

Track running through the project site.

On the advice of our manager in Sri Lanka, we broke the trip into two parts. Leaving the Colombo Hilton at 1pm on Tuesday we headed for Kandalama, planning to arrive at about 5pm. From there we overnighted at the Heritage Hotel and were picked up at 7am by a SLWCS four wheel drive for the remaining two hour drive to the site.

Brahminy Kite over the lake at Kandalama.

As we were driving we got to see first hand the 'conflict' areas between farmers and elephants, and had explained to us the traditional methods farmers use to cope with the problem. Namely, building a tree house on the edge of the jungle, and staying up all night to watch for elephants. Should an elephant, or more likely a herd of elephants come along, the idea is the farmer in the tree house makes a lot of noise to scare them off. Which doesn't always work. In which case the farmer can either choose to lose his whole crop, and income, or climb down the tree and beat the elephants off with a stick.

I would not like to be eye to eye with an elephant in the middle of the night.

I can't imagine the courage it would take to set to a group of wild elephants like that. I suppose if the alternative is total ruin, people can do some pretty amazing things. Even more amazing though, is that often it works, and the farmer can successfully defend his crop. Tragically, sometimes it doesn't work, and the farmer is killed or injured (50 people a year are killed during raids), and/or the elephants destroy the crop. At that point, things can turn nasty for the elephants, as the farmers family may start a vendetta and hunt and kill the elephants (estimated around 100 per year).

Humans and elephants have existed side by side in Sri Lanka for centuries, and farm raids were only very infrequent. What has changed in recent years, and made raids far more frequent, is three things;

- 1) Encroachment on native habitat by logging
 - 2) The spread of a type of non-native grass that has far lower nutritional value than the native varieties it has replaced
- Which both contribute to less food for elephants, and in some cases starvation, and

- 3) Farms planting cash crops of pineapple, corn and legumes - food elephants, particularly starving elephants, love.

Now, there are an estimated 4,500 wild elephants in Sri Lanka, but each year, the population declines by several hundred.

The SLWCS have initiated a number of innovative projects and trials, all with the same hallmark of improving or changing 'agricultural practices so that agricultural development can be sustained over the long-term'.

Hence we have the project we have become involved in - the Exetel Dairy Project - which, in a nutshell aims to:

* cross breed native and Indian cattle varieties to improve milk production from 4 litres per day to 10 - 15 litres per day. At this level, dairy farming becomes viable and can replace cash crops for farm income.

* Elephants and cattle co-exist, one eating the fodder the other leaves

* Sri Lanka imports 90% of its milk, improving local production will help the economy as well

* develop farming practices that minimize the impact to the local flora and fauna

* based on the outcome of the trial project, provide education to other farmers of sustainable dairy farming

* Increase productivity, which is decreased when the farmer has to guard crops all night

We arrived at the site by 9am and were greeted by typical Sri Lankan hospitality - a table set in the shade of a mud brick hut laden with local treats, and rich, strong Ceylon tea - thoroughly delicious and greatly appreciated.

Commencement of site preparation.

It being day 20 of month one of the two year project, there was not a lot yet to see at the site. Never the less, preparation work had commenced, and we walked the length and breadth of the boundaries talking with Chandee Corea (Operations Director) and Samantha Mirandu (Project Manager) who explained the project in more detail.

Part of the Project team: Chinthaka Weerasinghe, Samantha Mirandu, Darshana DeSeram, Chandee Corea.

What was abundantly clear was the enthusiasm and dedication that the project team has towards meeting the project goals. It is one thing to be told (albeit from a very trusted source) that this is probably the best way we can contribute to conservation in Sri Lanka, it is quite another to see first hand that there is no doubt it is.

Samantha Mirandu points out invasive non-native grasses contributing to the problem of elephant raids.

A couple of points of interest; as we were walking along the track at the site, Samantha pointed to a large round indentation in the ground, 'Elephant tracks' he said. 'How long ago do you think?' I asked. 'About half an hour'. We didn't see one, but it was pretty clear from those and other tracks, that, at any minute a wild elephant could emerge from the jungle. What do you do if that happens we asked? Run very fast, apparently.

Teak trees growing on the site form part of the long term viability plan. These trees are about 20 years old.

What we did see though, was wild pea fowl, a peacock with attending pea hens. Fascinating to see those birds in their natural state in the wild. Also, just briefly along the side of the road as we were driving was what looked like a rooster, but with a long, almost peacock like tail, and four, what looked like, pea hens - native jungle fowl as it turned out.

Posted by Steve Waddington at 13:07

Thursday, December 17, 2009

PeerApp Commissioning Today

Didn't work out as we had planned. Everything was full prepared for the cut over at 3am this morning, but, the best laid schemes etc.

What went wrong (Wha' gang off' agley? It is nearly the time whatever is misquoted from Robbie Burns' original words will be more than made up for by inebriation, after all) in this case was two things:

1. Of the four GBIC interfaces needed to carry traffic between the cache and the border routers, one, which worked fine in testing yesterday, wouldn't establish a connection. Nor would the spare. So we had to figure it was a port problem with the port itself.
2. All four ports on the cross-connect switch reported zero bytes per second throughput, even though the interface counters were incrementing bytes in and bytes out. Not a fatal problem in itself, but still, an indication of perhaps some other problem or even IOS bug.

So those two issues ate up the maintenance window time before they could be resolved and that was that for the morning.

Next step is to retest and fix whatever is causing (1), and fully understand/fix the reason for (2). Which means the cache won't be commissioned this week, and most likely not until at least next Thursday. Although, it may even be better to wait until the period between Christmas and new year, which is the lowest usage period of the year.

We will have our weekly network and bandwidth review meeting tomorrow and decide then.

Posted by Steve Waddington at 08:56

Monday, December 14, 2009

New PeerApp Installation

Most of the new server equipment arrived on Friday, the remainder coming today, along with the PeerApp engineer who will be working with us on the installation.

All of the equipment is now at the POP, with the estimate being that it will take another full day to complete the physical installation, and remove the old equipment. If that work out, then the maintenance window from 3am to 6am Wednesday morning is our target timeframe to commence sending traffic requests to the new cache.

One complication that may delay things however, is that we have a pending 'statutory carrier obligation' request which has to take priority over other work. If that work turns out to be more complicated than we expect, then we will have no choice but to see it through, by necessity holding up the cache setup completion.

Worst case, I would say that we should be able to see early results by the end of the week - or perhaps over the weekend, depending on how long it takes the cache to populate.

Posted by Steve Waddington at 16:04

Friday, December 11, 2009

1Gbps to Every Home in Australia..

.. well the ones that matter anyway. Which I take to be the oft quoted '90%' of Australian households.

While Conroy is going on about the brave new world of an NBN, and how if the Luddite detractors had their way, the 'government' would not have built an electricity grid. (It was J.P Morgan, Westinghouse and General Electric in the US that paved the way, but lets just suspend historical fact for the sake of a political sound bite). Why is he being so defensive if it is such a self evidently good thing?

The spin has clearly changed from 'Everybody wants this' to 'it is up to the government to do this for the good of all (even though they may not think it is such a good idea now)'.

Let's put aside that 100Mbps+ bandwidth is already being delivered to many households, and many more could have it if they wanted. That upgrading existing infrastructure to accommodate bi-direction, true IP, traffic can be done with the upgrade of a few boxes (ok, a few thousand boxes, but for well less than \$1bn in any event).

The thought occurred to me, why settle for 100Mbps? Whether it is fibre to the node, or fibre to the house, where does the magical figure of 100Mbps come from? 1Gbps speeds have been common, and cheap, for half a decade now, on both copper and fibre.

So what would it cost to deliver 1Gbps to every (90%) house in Australia? I did a few 'back of the napkin' sums, and, using current retail (or street) pricing for equipment, my reckoning comes to a shade over eight billion dollars.

That is for fibre to a media converter to each house (8,000,000), terminating at each exchange (1,100) on a 1Gbps switch port, aggregated into a suitably large router for backhaul either to a central point in each state, or to the call collection area (66) for the region and then to a central location for the state.

To calculate the termination cost at each house, I used the current price Telstra would charge for a new service, with a mark up to allow for the extra termination cost of fibre.

Balanced against all the costs I have trivially over looked, has to be the fact that I based all the capital equipment costs on the retail price. One would have to assume some sort of discount would be available for an order of eight million. Also, that a contractor terminating every house on every street is going to charge less than the once off special truck roll the Telstra pricing for installation reflects.

But, \$8bn is still a lot.

Actually, no it isn't. On a 5 year ROI, it is \$17.03 per month per household.

1Gbps to my house for around \$20 per month. I think I can see the value in that.

Who has the failure of vision now Mr Conroy?

Posted by Steve Waddington at 14:28

Thursday, December 10, 2009

The times, they have a-change'd

(apologies to Bob Dylan)

This year, more than any other, I have noticed a change become apparent. And by a 'change' I am a fundamental shift in the paradigm of business in the service provider industry.

I am never quick to pick these things up, so I can only assume that if I notice a change, it has well and truly already happened, and many other people are enacting/changing/making their plans accordingly. Maybe they are all keeping quiet about it, or maybe they are in denial, but from what I see, there is little evidence of accommodation for the change actually being made. Maybe it is something that will be more noticeable after Christmas.

What is this mysterious 'change' then?

Perhaps it is best explained by its ultimate effect - which is not yet, and there is still quite a while, before it is fully realized. So here it is:

"There will be no job in Australia for anyone who is not directly contributing more to the GP of their company than their employment cost"

But surely that is already the case? No efficient company employs anyone if there is not a commercial reason to do so - at least in theory. Which I am sure is correct, if not always in fact, then at least in desire of the people controlling the finances.

However, how many people go to work each morning thinking 'Today, the company I work for is going to pay me 'x' hundred dollars, so what will I do today to ensure the money is there to pay me'?

I would take a bet that a far more common attitude is 'Work today, only 'x' days to go until the weekend', followed by whatever minimal effort required to do what they are instructed to do, while attending to the important things like Facebook and MSN.

Which has been the way for decades, at least, of employees doing the minimum required on the expectation that the 'big rich fat cat company' they work for will keep paying them (them or someone else, but better the devil you know), and employers accepting that.

I am not arguing that it is right or wrong, or that it is true in every situation. I am saying that it is my observation, from what I have seen, people I have spoken to and heard, over many years.

I put forward the view that this attitude/acceptance is, if not the, then a major driver for for many companies outsourcing to countries where wages are much lower than Australian award wages. If as an employer, you have resigned yourself to pay staff to do a minimal job, would you rather pay them \$1,000 per month or \$5,000 per month? A no brainer, which is just as well, given the very low amount of brains that led to that situation in the first place.

For a brief time I worked as a consultant for a 'big five' IT consulting firm. It was very clear, as a senior manager, just what was needed to stay employed there. Time accounted for in six minute blocks, 80% of time billed, and get as many of your colleagues as much billable time on any job as well. Every business run as a practice knows those figures - without billable hours paying for staff, you just can't have the staff.

Why is it so much less 'cut and dried' for a service provider? Is it really such a hard exercise in abstract thought to see that the time a network engineer/sysadmin/programmer/support person spends doing things can't relate directly to the bottom line of the business. And if that major hurdle of thinking can be overcome, then surely it is a small step to a) relate it to real contributing value to the business and b) work out how to ensure you are valuable enough to overcome the competition (being the people who will work, very happily, for 1/5th of your salary).

What I can say is that I can define a single point as the big change in my career from being 'just another engineer' to someone who is highly valued, given the most important projects and promoted, by whatever company I worked for. And

that was, one day (I am not sure of what exact time period the thought crystallized, it was over 15 years ago) I embraced this thought "Whatever the company I work for pays me, I will try to ensure I return 10 times that in dollar value". Which is a very simple way of putting it, so I will qualify it a bit:

- if I was on, say \$50k pa salary as a pre-sales engineer, I would ensure the sales I contributed to exceeded \$500k pa
- If I was responsible for CAPEX for a project, I would ensure the ROI was as short as possible, and look for enough projects over a year to return more than ten times my salary
- when first starting a new job, I would look for an immediate cost reduction (which I never had any trouble finding) that would at least cover the cost of my starting salary

As a junior engineer, there was little opportunity to have much control over what work/projects etc my employer of the day could offer. But the funny thing is, once I had embraced that mode of thinking, from that point on, I did not stay a junior engineer for very long.

Posted by Steve Waddington at 10:55

Friday, December 4, 2009

The problem with Google's Public DNS

At first look, two problems I see are:

* how will it handle captive portal systems?

These are an increasing part of ISP operations, redirecting end user traffic for any number of reasons. End users switching to the Google public DNS servers will cause portal redirects to break, most likely resulting in a hung browser session or 'page not found' error

* how will it handle internal IP address maps for private networks?

Some ISP's, and many corporate networks make extensive use of RFC 1918 IP address space. Company intranets are a prime example. Mapping that through a local DNS resolver is no problem. Completely breaks with a public resolver though.

It is interesting to see a selling point from Google is that it will "Get the results you expect with absolutely no redirection.". I assume, one day, Google plan to make some revenue from this system. I wonder how they plan to do that?

Posted by Steve Waddington at 12:51

Tuesday, December 1, 2009

What does the NBN mean to a real broadband user?

I was driving with my oldest child yesterday, who has just completed his first year of Uni. As a teenager, young adult, FPS and MMO gamer and student, it is hard to imagine any demographic that a) would make more use of the Internet in their day to day lives, and b) would have a better lay persons understanding of the Internet and broadband.

So he was telling me that the consensus amongst his friends was that the much high speeds provided by fibre would mean lower game pings. Taking a typical Call of Duty in game ping time from 80ms down to 20-30ms.

I found that comment very interesting for a couple of reasons.

First, I asked him 'What about download speed?'

To which I was pleasantly surprised to hear that download speed depends more on the speed of the server being downloaded from, 'and anyway, we have 12Mbps don't we, so stuff downloads about as fast as it can?'

Further questioning revealed the first part of his comment came, he thought, from something I had told him at some time (so kids do listen to some things we tell them, wow). The second part was from comparisons he and his friends had made based on their various broadband connections and service providers, and that a uni tutor had told him the University network limited external connections to 4Mbps per user.

Download speed of current copper delivered broadband doesn't seem to be an issue - at least for one particular group of, semi well informed, users.

Next I tackled 'how will a new NBN improve your ping times then?'

'Well' I was told 'In the US everyone has T1 or T2 connections and their ping times are 3-15 ms.'

Uh huh.

In this context, 'T1 or T2' was meant to mean 'fast fibre connections to the home', so for the sake of maintaining focus on the core issue, I let it go at that.

I went on to explain that ping time is a function of user distance from the server, and once the line speed went over about 256kbps, faster speed made very little difference to the ping time (which we all know, of course). I likened it to a freeway (since we were driving down one at the time), whether it was two lanes or four lanes or twenty lanes, the speed limit was still 100kmph.

'But can't the data go faster than the speed limit if there is more room?'

'No son, the laws of physics police that one' I informed him in sagely father tones.

It required a little more explanation than that, but the concept was quickly enough understood.

Never the less, the question remained that wouldn't it still be better to have 100Mbps Internet than 12Mbps?

'Possibly' I said ' but what would you use it for?'

Receiving the stock answer of 'movies, TV and streaming multi media'.

So I went on to explain that we already had a connection of well over 300Mbps bandwidth when we had Foxtel. And now with free to air digital TV we still have well over 100Mbps of bandwidth for the express purpose of TV and movies. Foxtel IQ and Tivo filling in the 'time shift' gap, if there was one, left by the demise of the VCR.

I am not sure I fully got the concept across that broadcast media is just the same as Internet bandwidth when there is no, or minimal, need for real time interaction from the user. However, certainly for the time we were driving, we couldn't

come up with a valid example where it would not be equivalent (having excused P2P downloads as something to be addressed by the content providers in future - that is, if the content is readily available to meet consumer demand, P2P of pirated TV and movies will become irrelevant).

The conclusion we arrive at therefore was the _desire_ for faster broadband for the informed residential user was literally, to make the ping times faster rather than to have access to more bandwidth. But, as we see, that is a delusional wish based on misunderstanding the physics. Assuming \$43bn could ever be justified for a 20ms ping reduction in Call of Duty 4, of course.

Posted by Steve Waddington at 10:50

Friday, November 27, 2009

New PeerApp for the New Year

After quite a bit of discussion, checking and double checking, we have reached an agreement with PeerApp to go ahead with the upgraded platform in Sydney and also add new equipment for Melbourne and Brisbane (since our traffic is now at levels in those states that make caching worthwhile).

The equipment will be the latest iteration of the Ultraband product we have been using, which upgrades:

- bandwidth handling capacity in Sydney
- the ability to better recognize cachable protocols, and improvement to the traffic redirection to the cache engine
- allowing the inclusion of multimedia and other web content

We are expecting the equipment for Sydney to arrive and be installed mid December, with the goal to have the equipment in trial phase before Christmas. Melbourne and Brisbane installations will subsequently take place later in January.

With the upgrades to the equipment, plus increased disk capacity, we are expecting cache generated traffic to be in the order of 600Mbps in NSW, 250Mbps in Victoria and 150Mbps in Queensland.

Posted by Steve Waddington at 14:55

Monday, November 16, 2009

HSPA For Fleets

For organizations with many people in the field using HSPA, we can assign all IP addresses to those services from a single IP address block.

Alternatively, we can assign each service from a private IP address block. This may be an advantage where the client wants to strictly control Internet access for all users or restrict access via a central firewall only.

Posted by Steve Waddington at 10:48

Monday, November 9, 2009

hackers - An excuse for incompetent IT management?

Has anyone thought about who or what a 'hacker' actually is? Apart from the root of all evil (in IT), a justification for a huge security budget, and like every bogey man before them, a way to scare children allow hacked phrases to pass for policy by those too lazy to spare a moments thought on the subject.

I saw this report, and it just made me wonder; instead of reporting "Power outage caused by hackers", would it not be more accurate to say "Power outage caused by IT idiot unable to secure vital infrastructure systems from teenagers"?

No, no - far better to blame the outage on these mysterious and vapourific 'hackers'. Because everyone knows what a hacker is, how inherently evil they are, and a two day outage is just the sort of mitschiefe they would cause. For no other reason than they can. Whoever 'they' actually are.

It is very convenient that invoking the word hacker seems to completely negate the need for any further investigation into the cause. Everyone knows you can't do anything about hackers, right? And I am pretty sure by now the logs will show whatever they need to show to support the 'hacker' case.

So lets just say it was these scallywag hackers for a moment. Who are they meant to be exactly?

No one extorting money, or doing it for any other reason than vandalism it seems.

That would seem to put it in the realm of teenagers with a pretty good Internet connection and too much time on their hands. And we can all agree there must be lots of people like that.

Exclude the ones who idle away their time with WoW/Quake/Halo etc. Which must really narrow down the field a lot. Exclude those that are not inherent vandals (my kids and yours, at least). Then exclude those that are vandals, but express it with a spray can or a knife on a train seat.

What does that leave us with? Someone who wants to create widespread vandalism by, specifically, hacking into a public utility, with (the real kicker) the ABILITY TO DO SO.

Because of course, it is pretty easy to hack into a computer system, right?

Wrong.

It may be easy enough to hack into a pirated, unpatched, ten year old Windows OS running on a computer in China or Romania. I think however that the same is not as true for fully maintained servers run under the supervision of a competent IT person. Even without a zillion dollars on firewalls and other security enhancements, just the OS itself - Windows, Unix, Linux or whatever, is enough when run by someone reasonably competent.

But even forget about the inherent security in OS's. Look at it another way; saying a problem is caused by 'hackers' is the same as saying that a teenage vandal knows more, can out think, and can beat the experience of a (presumably) degree qualified IT professional who is paid by their company to secure against just such an event.

Well, any number of US TV shows and Hollywood movies show us just how common teenage geniuses are who can out think any corporate 'suit'.

Or, perhaps the true headline should read "Power outage caused by ass covering IT idiot, tries to blame it on hackers".

Posted by Steve Waddington at 18:00

Friday, November 6, 2009

Customer Feedback Forms

I think I fill out more than my share, because I don't mind doing it - on the assumption that whatever company has gone to the trouble of asking for it will probably, at least in some way, act to improve their service in some way.

Over the years, the most consistent, to my observation most acted on, and in my opinion, best put together, are the feedback forms requested by Cisco at the closure of every TAC case. In the past I have had direct calls from senior managers in the Australia and the US in response to - not even particularly - negative evaluations I have given. Leaving no doubt there was a genuine desire in the company to fix whatever it was that led to my dissatisfaction in that case.

And, as many people would know, they really hassle you to fill out the form too. The last time I was in that loop it was three emails, of increasing urgency, along the lines of 'We haven't received your feedback yet, please do it now'

I think Cisco used to say that any customer response of '1' (the lowest response on the 1-5 scale) would be reviewed personally by John Chambers. I don't know if that is still the case now, but nothing led me to think it didn't actually happen that way. Though smaller than they are today, it was still an amazing thing to see in, what was then, a \$15bn pa revenue company.

I don't know if Cisco set a new 'industry standard' for customer feedback in that regard, but certainly many companies in many different industries now have a similar approach. From just about every hotel I have stayed at in the last two years, to airlines, and even the dealership where I get my car serviced. Similar sorts of questions, similar sort of feedback and follow-up.

But here is the kicker - they ALL make one fundamental error (only my opinion of course) in the way their survey asks for information - Where the ask, and most of them do, a question about 'your expectation'. It will usually look something like this:

On a scale of 1-5, one being failed dismaly and five being far exceeded, how did our service X meet your expectation?

And there will be a bunch of questions about different areas of the product or service where the respondee is expected to give a rating on their 'expectation'.

One can only assume that somewhere along the line, those customer feedback forms contribute to or form a basis for and employee/manager/department evaluation. And fair enough too.

But nowhere I have seen, on any form, have I been asked to actually say, from the start, what my expectation was.

So what am I actually being asked to say? My assumption may be quite wrong, but the feeling I get (and the first hand information I was told by some TAC engineers many years ago), is that the goal or KPI is to get across the board '5's' for every response. Which means the the company has to 'far exceed' my expectations.

No problem, always nice to know that people have an incentive to exceed expectations.

But what happens if my 'expectation' was already for excellent service? And I received excellent service. And therefore to answer the question as it is asked, I really need to put a '3' - the service met my expectation.

It seems like a very big error in the data collection method to me. Maybe I have missed something fundamental in the way feedback is intended to be analyzed or some other subtly that completely escapes me. It wouldn't be the first time. Or maybe, the question/form structure is just being used by rote by company after company, with no real understanding of the true intent of the person who was the original source of such things.

An example of this is where last month I stayed at Star City while I was in Sydney. A really nice, top class, 5 star hotel (A real bargain at the off-season room rate I got on-line). My expectation was that the service would be great, the room would be first rate and I would enjoy staying there. And it was no surprise to find my expectation exactly met. Which is what I wrote on the evaluation form.

The fairly prompt email feedback from the hotel was 'We are sorry to hear you did not enjoy your stay'. But I did! My expectation was met in each case - the middle response on each question they asked. What I think they, and every customer feedback for like it, should do, is at the very start, ask what my expectation actually is.

True, it will mean the analysis of the response will be more complicated than just averaging the gross total. But surely that is what is needed if the intent of such feedback is to be fulfilled.

Anyway, the result of my musings on that subject are that I have redesigned the corporate support feedback page to this:

Please select a response from 1-5 for each question using the following guide:

- 1 - strongly disagree or far below my expectation
- 2 - somewhat disagree or below my expectation
- 3 - no opinion either way or just met my expectation
- 4 - somewhat agree or exceeded my expectation
- 5 - strongly agree or far exceeded my expectation

If you know the ticket number, please enter it here:

If you know the name of the person who handled your fault, please enter it here:

1. Before I logged the fault, I viewed Exetel as a high quality service provider
2. My fault was handled promptly and efficiently
3. The fault was fixed in the timeframe I expected
4. The fault has been resolved to my complete satisfaction
5. I was kept fully informed of progress at every stage
6. I think Exetel did everything required to ensure the fastest resolution
7. The way this fault was handled has raised my opinion of Exetel

And lastly, the other minor frustration I have with such feedback forms is, there is invariably no way to say 'you haven't asked me the right questions!'.

Hence; today's blog post.

Posted by Steve Waddington at 13:04

Wednesday, November 4, 2009

How to put an end to illegal P2P file sharing

Now this is just a thought. But would it be possible to turn the greed of copyright infringers against themselves and use that, where enforcement (or lack thereof) methods have failed?

Think about this; P2P file sharing can only work if enough of the people who download the files also make them available for uploading - ie sharing - as well.

What if everyone only downloaded files, and there was no, or very little resource dedicated to uploading, and resharing the file? I think in a very short time, the novelty of getting the latest movie or TV program 'for free' would wear off, particularly if that programs download time were to be measured in months rather than days.

How much P2P copyright infringement went on when everyone only had dial-up speed? (well plenty apparently, but that was music files @ 3Mbyte size, not movie files @ 700Mbyte size). Even so, mp3 sharing would not have worked if it wasn't for the 10:1 compression of the music format in the first place, and the relative speed of 56k modems compared to, say, 9.6k modems.

The question is, who is going to wait a month or more to get a movie file? A couple of days maybe. But even then, those couple of days better not cross \$1 Tuesday at the local video store - which will represent a much better return for time/effort/'I want it now'.

Ok, so a scarcity of upload resource which consequently leads to very slow downloads may well put an end to illegal sharing. But how do you convince all the 'community minded' (it cost me nothing to make, or get, so why don't I make it available to you as well) sharers (thieves) to not on-share what they have downloaded?

I can imagine a straightforward way to do just that.

Make it known that there is no penalty to download - only to upload. That is; a campaign of clear public statements that downloaders are simply 'victims' of unscrupulous people who are preying on them, and banking on them, to propagate their crime. The downloading of a copyright file is not a crime - or if it is, will not be pursued, only the re-sharing and uploading of it will be.

And inform everyone, via the many, many public forums available to do that, of just how to set their P2P clients to 0 or 1kbps upload, while maintaining full download speeds.

It seems to me that would accomplish a number of things:

1. Illegal P2P would die a natural death in a few months, and the sites like Piratebay that profit from that
2. Legitimate P2P would be unaffected
3. It would have to reduce the vilification the recording and movie industry comes in for for targeting home users
4. It would make it much easier to identify the people that are making commercial gain from the copyright theft, because they would be the only IP addresses still seeding

Anyway, just a thought.

Posted by Steve Waddington at 19:35

Wednesday, October 14, 2009

Customer Support

I get all the emails, as do the other directors, send via the suggestion box on the Exetel user web page. It is really great - such a basic idea to gain feedback - many good suggestions, literally dozens of which we have implemented.

Yet, after genuine suggestions, and some people using it for reporting faults, the next most common (about 3 in every 100) 'suggestion' can be paraphrased as 'Improve your customer support'.

My reply in each case (yet to be answered by anyone) is always:

Dear Sir (or Madam),

I am sorry to hear we have not met your expectation in this case. Can you please let me know the problem you are having, the ticket number or any other information and I will investigate the case.

In another life, many years ago, I would occasionally get a call passed to me (meaning many people had already tried to deal with it) that would go something like:

"Your customer support is terrible"

"No Madam (or Sir), it can't be..." I would start to say

"Well I am telling you it is, and you are apparently a senior manager so what are you going to do about it?"

"As I was about to say Madam (or Sir), our customer support can't be terrible, because we don't offer any." and in the slight pause of stunned silence I would add "but if you want to tell me what your problem is, I will see what I can do to help you."

Ah, Windows 95 and dial-up modems. Those were the days. Admittedly not very good days as far as home user ease of setup for Internet access, but days never the less.

The philosophy then (abandoned after a short period of time as I recall), was that there were so many variables, and so much complexity (it was a dream is a customer had Windows 95, because in those days it was not uncommon for people to be using MS DOS), that it would be impossible to provide over the phone support that would actually help anyone anyway. So just don't do it, and the people that can set themselves up will, the rest will do the sensible thing and get their local computer guy to do it for them.

But, as every other idiot service provider started to offer 'The best customer support' as part of their marketing - almost completely ineffectual in helping anyone, other than to either read out the instructions over the phone the end user should have read themselves in the first place, or, eventually, advise them to seek the aid of their local computer guy - then so too did we.

Something I probably considered a loss of nerve at the time, but undoubtedly inevitable. I guess.

Anyway, what does that have to do with anything now?

I was just thinking, with unjustified nostalgia, about our initial premise of customer support at Exetel. Very simple - Deliver a perfect network, and there will be no need for customer support. Or as John would often articulate it 'If the light on the modem is not flashing, there is nothing Exetel can do to help with whatever problem is being reported'.

But what I liked the most was the actual definition of how we defined 'Customer Support'. Which was (and still is); The best support is to provide the best service in the first place by making it as reliable and available as possible, provisioned as quickly as possible, and when broken, fixed as fast as possible. And, when someone has a problem, give them the best help we can to get it resolved as quickly as it can be.

So we hired, and hire, engineers with good degrees, people with knowledge and people with passion for the technology.

To my mind that is still the best possible customer support that can be offered - because it has the best chance of the fastest fix to any problem. Which just has to be better than an Irish backpacker in a generic call centre, right? Or whoever it is large call centres are staffed with now days.

Sure, there are people who are simply vexatious and will, given the chance, argue for the sake of someone to argue with in their otherwise uninteresting lives. Thank god we have the TIO who will lend an ear to anyone who decides to call them for any reason whatsoever, charge the ISP for taking the call, and should the ISP dare to dismiss the frivolous, and sometimes outright lies, of the complaint, well, straight to 'level 2' and another \$400 charge to the ISP. But there is always the tail end on any bell curve, so there is not a lot anyone can do about that.

Perhaps the fact is, if religion teaches us anything, that it even something perfect (God, we are told, being a perfect and omnipotent being, His word can be no less) is open to many different interpretations and/or misinterpretations, depending on an individual's view point.

Still, in my own very small and humble way, I really would like to know why the three in one hundred people who take the time to write a suggestion think our customer support is less than they were expecting.

Posted by Steve Waddington at 10:08

Tuesday, September 22, 2009

Cloud Cuckoo Land of Communications

Just read this.

I see. The 'industry' is the one to blame for the debacle. that is the lack of communications strategy the government has so far come up with. Not that I have any particular view that the 'industry' can effectively self regulate - but apart from the one monster monopoly hogging most of the pie and playing spoiler to its competitors, the 'industry' has done pretty well offering competitive services on a national basis if you ask me (in which, I immodestly like to believe Exetel has played no small part).

Fortunately, the rest of the statements Conroy is quoted saying only go to show just how lacking in substance, meaning, or reality anything he is likely to say is.

Satellites. ooohh aaahhh. Now we are talking real technology bring to bear on the problem (there is a problem right?). Wont that make a great press shot Conroy and Rudd standing in Mission Control Centre (someone pick a marginal electorate were that will be located), with a Really Big Screen in the background.

Of course, with all the impressive technology flying around the sky, who will even notice the 300ms latency added to every packet, and the destruction of VoIP and other real time services it will cause. I guess Conroy will just have to say "We are fed up with scientists telling us what the speed of light is. If we can't do a deal with the laws of nature, we'll make up our own reality"

Somehow, I think they already have.

I mean, fair dinkum (mate). Do you know what banks charge for ATM fees? Have you seen their profit reports lately? If I made an average of 4% on every cash withdrawal at the cost of a couple of hundred bytes of information transfer, do you think ANYTHING would stop me putting an ATM anywhere where money could be spent?

But WTF (if Conroy doesn't know what that means, I am sure his boss can fill him in) does that have to do with residential broadband Internet? Hmmm, I see every service station sells mars bars at the same price, so petrol should be the same price everywhere too. And it would be the transport industry to blame for that debacle, according to Conroy's logic.

Not a good analogy of course, because petrol has to be trucked by real trucks driven by real people who have to house and feed their real families. That harsh fact adds cents per litre. (so what, the Mars bars just walk there by themselves? Ok, ok, I said it wasn't a good analogy, you have to consider transport cost per unit as a percentage of the total cost etc.. Suffice to say, Conroy comparing the massively profitable service of ATM's with the minute margins available on Internet services is not valid)

For IP data however, the traffic is over an already in place network, paid for (many times over) by telephone call tariffs. Sure, there are upgrades and maintenance to pay. But what is the real cost of that? yes, there are a set of figures provided by the incumbent monopoly. They seem a little high to me, I wonder why?

Posted by Steve Waddington at 14:57

Monday, September 21, 2009

A very pleasant prospect

We took the opportunity last weekend to visit an 'aged P' in Albany. A four hour drive down the Albany Highway from Perth, but fairly relaxing in the cool late winter. And nice to see the green and pleasant hills roll by through Williams to Mount Barker - auguring for a good wheat harvest.

With the HSPA stick plugged into her laptop, my travel companion reported a usable signal (to update facebook so the kids would know we had made it that far at least) as we approached and passed through Williams, out to about 10 k's either side of the town.

So too in Kojonup where we stopped for a mid trip treat of a pie from the local baker. No problem to check email and reply to a couple of forum posts from the cafe. Download speed reported by speedtest.net was 2.1Mbps.

And so, aged P visited, a nice dinner with old friends, sleep in on Sunday. Now what to do for Sunday lunch before heading back home in the afternoon?

Denmark, pretty any time of year, but particularly beautiful on a crystal clear south coast winters day.

The latest James Halliday guide in hand we stopped first at West Cape Howe wines (sadly sold out of the award winning Cab Sav), and were directed from there to Forrest Hill winery in response to our question about a nice place to have lunch.

Four km from Denmark as the crow flies, overlooking two verdant valleys and a little stream, with glimpses of the Wilson inlet through the interposing hills. Just how more pleasant an relaxing could you get? Now, I completely understand the view that many people would have that it is just not appropriate, perhaps even verging on sacrilegious, to, at this point, take out the laptop and check on the business.

'Work' at this time and place? Just how big a philistine can you be?

I suppose if you view 'work' as 'work' and a chore that is an inconvenience to your far more important personal and social life, then that is indeed a valid point of view.

On the other hand, what better time to make a reasoned and unhurried response to a suppliers proposal, a customers suggestion, or a colleagues email? I am sure the people receiving the emails appreciate it more, in any event.

Even Zeiss optics at 5 megapixels don't do justice to the Forrest Hill Restaurant nestled in Hills of Denmark.

Four kilometres from the tiny southern town of Denmark. Ancient forests intermingled with green pasture and vinyards. And near perfect HSPA reception. 'Almost anywhere' is not a phrase that Telstra like anyone to use. I suppose when, as a commercial company, you have the windfall of a network paid for by the previous 90 years of being a monopoly, the massive advantage of the national coverage that provides is not something you would want to let go without a (dirty?) fight.

Never the less, all I can say is that in my travels around the countryside, just about anywhere I go, where I would reasonably expect to have HSPA access, it is there. But also in many places where it is unexpected.

That aside though, being able to go almost anywhere (a reasonable person might want to go) and not be bound to a wire line connection - what a very, very pleasant way to work.

Posted by Steve Waddington at 11:12

Thursday, September 17, 2009

Malware affected PC's

Apparently a new code of conduct has been drafted by the IIA to disconnect customers that have infected computers.

What a great initiative, and you can see the ISP's themselves are totally behind if this quote (from the itnews article) from Steve Dalby of iinet is indicative:

"Potentially it's something that we would do. If there were some costs we might consider whether government funding was available, but again it's very hypothetical,"

Now don't over commit your company there Mr Dalby.

How 'potentially' and 'hypothetically' wimpy is that and how sad to put a hand out for government funding. So is that what is required - Government funding - to be a good net citizen is it?

It seems we, ISP's as an industry, need to government to fund anything we might do that is decent or right - filtering child porn, stopping malware spam. Have we come so moribund and incapable of innovation that anything other than connecting a customer (which we can mostly do, most of the time) is now beyond us?

As for costs. Ha. When Exetel started doing four years ago what the IIA has now just drafted, our development cost was a few weeks of programmer/sysadmin time (in conjunction with their other work). I don't know, call it \$2,000 maybe.

The running cost for that automated system is nothing, or as close to as makes no difference. The actual cost is in fact a cost saving, because the automated system handles almost all of the issues without tying up staff time - my rough estimate would be \$50,000 saving or more over the years.

Anyone who thinks implementing such a thing would cost a lot must either have a poor grasp on the concept of automated systems, or very, very bad advice/ability from/in their development team.

But the true shame is the need for a code in the first place for what is simply a responsible practice that should be done as a matter of course.

Posted by Steve Waddington at 09:54

Monday, September 14, 2009

An end to P2P caching?

While the traffic generated by our P2P cache - the PeerApp UB2000 installed almost two years ago - has remained the same, our bandwidth has increased quite a bit over that time. Which is to say, the ratio of traffic delivered from the cache vs external bandwidth has decreased over time.

Reasons why that may be that we have considered are:

- Less P2P traffic (ha)
- More use of protocols such as Rapidshare that use port 80
- use of encryption within P2P
- more local trackers in use by end users so more traffic via peering

My guess (or rather my lead engineers guess, which is better informed than mine), is that the second and last point are the major contributing reasons.

Never the less the cache generated traffic has represented a considerable cost saving over the time we have used it. And that has been one of the contributing factors to keeping us keenly priced in the competitive ISP market.

With decreasing wholesale bandwidth pricing to Australia as well as the smaller portion of traffic the cache accounts for, right now the price of generated bandwidth is only slightly lower than direct bandwidth. And with offers from some suppliers, the price of direct bandwidth will drop bellow the cache generated cost.

So, we will have to see what can be done on both the cost front and the traffic generation front. One thing we will trial in the next few weeks is port 80 only caching, which, I have been told, now accounts for a lot more general web traffic than P2P (because of Rapidshare? I don't know).

It will be interesting to see how that turns out.

Posted by Steve Waddington at 10:44

Wednesday, September 9, 2009

The Best NOC...

... is no NOC.

Gasp, shock, horror! But indeed it is true, at least in terms of the 'NASA Ground Control' type NOC, with the Really Big Screens typically envisaged when we think about a Network Operations Centre.

It was possibly true that in the 50's and 60's, complex globe spanning operations could only be coordinated from a central location where all the monitoring technology could be concentrated. And possibly also true, that only a technician sitting in front of an array of screens and buttons, next to lots of other technicians doing the same (and with one or more Really Big Screens at the front, should they ever look up), could take the corrective action needed if something were to go wrong.

Yet having seem may in current operation in carriers and ISP's alike, my view is they are just a waste of time and money. Serving no useful purpose that could not be just as well, or better, be done from anywhere with simple web interfaces.

The reason those centralized faculties may have been needed in the past is, I think, because with the communications systems of the day (think dictation, typing pool, inter-office memo and mail delivered by a postman) that was the only way information needed to run any complex system could be passed on at the speed it was needed. Which was; to be sitting next to, or a short walk from, the person that needed to have it.

(Except of course for by telephone, but it is a bit hard to convey the details in a technical diagram over the phone. (Except of course for fax, but is it a bit hard to coordinate faxed and verbal information that needs to be shared with many people. (Except of course for conference calls. (But I regress too much))))).

Suffice to say, that email, IM and all the good things we have in this age of communication mean that very, very complex operations can be managed without any centralized facility. That, and script languages that allow easy integration of back end systems with notification systems and web pages.

For a start, what is the point of monitoring and those Really Big Screens? Presumably, a red light flashes on the screen indicating a problem, a supervisor notices (the whoop-whoop siren might also give him a clue in the case of a major problem), and some technicians frantically press one or more of the mysterious but impressive lit buttons on their customized consoles to fix the problem (or at least silence the whoop-whoop alarm). What drama and excitement. At the end of the shift, everyone can go home satisfied at a job well done.

Ok, replace all that with a script or two that, when the monitoring system would cause a red light on the big screen, actually does something useful like page or sms an engineer who can fix the problem. Much less drama, and no where near as much fun, but, just a bit more efficient.

Does the engineer need to be at Mission Control Centre? Maybe, if they are the sort of engineer that can't be trusted and has to be watched over. Just the sort of person you want working on your critical systems in a crisis, right?

So let's assume we have hired reasonable, competent engineers who actually don't need sweat shop style management practices to be able to do their job. Since we have global communications and messaging really is effectively instant to anywhere on the globe, it makes very little difference where the engineer is when he gets the alarm, as long as there is access to the super customized, proprietary mega console so he can press the right flashing button to fix the problem.

Or, we could just link the routine operations to a web page that (again using one of those very handy script languages) lets the engineer access most things he needs with a click on a hyperlink. And there is always telnet and ssh when the going get really tough.

'No NOC' is how the Exetel network is managed anyway. Everything that can be is on a web page, and most things that can are automated with scripts. We all have mobile phones, and either HSPA or ADSL, or both. Scripts and other systems monitor everything we can think to monitor every minute, and any problem is messaged out to one or more people who can deal with it within seconds of it being detected. Even the escalation process is automated, so if the

problem isn't attended to by the duty engineer in the allotted time (either 5 or 10 minutes, depending on the problem), the next message goes to the manager responsible (and ultimately to me if nothing is done by anyone for longer than half an hour - pleased to say that level of escalation has never been reached)

I would like to say the methods were all developed here and refined over the last five years. But in fact, when I was thinking about it to write this blog, I realized that never, from TPG in 1995 on, on any of the national and International networks I have build and managed I have never actually set up a centralized NOC. There just was never the need, for all the reasons above.

All I can perhaps say is that the methods have become more refined, and more cost effective at Exetel. And I think that is because of the permanent, persistent communications connectivity we all have (which was not the case in the 90's with mostly Dial up Internet), through the ubiquity of mobile phones, home ADSL and most recently HSPA.

Posted by Steve Waddington at 14:43

Tuesday, September 8, 2009

Businesses run on ADSL

Residential grade ADSL at that. Was one of the more surprising things that cropped up during the recent IP address change we undertook for some 10,000 customers.

Most of them I doubt even noticed, or if they did notice, care. But at least a dozen end users (and maybe a lot more) had concerns like these:

- hosting a number of servers
- running several web sites
- terminating VPN's for remote users
- hosting email for several different domains

There are two things that struck me about that. First, from an 'Internet access is critical for my business' aspect, cheapskating by using residential ADSL, and then complaining when something changes really is not the best way to go.

But on the other hand, it is amazing to see that residential ADSL has become so reliable, that people can actually (and do) use it for, critical, business applications. Even though there is no guarantee of continuity of service, the actual availability is such that it is taken for granted for those types of use.

Posted by Steve Waddington at 09:54

Thursday, September 3, 2009

The best IT department...

... is no IT department.

Which is the basis on which Exetel has operated since we started in 2004. Which, as a startup, with our first few staff, of course made perfect sense.

There is a very funny series of sketches from the Saturday Night Live show that feature 'Nick Burns - Your Company's Computer Guy', where Jimmy Fallon as 'Nick Burns' berates and ridicules the computer users in the sketch who 'knows less than he does' - which is of course everyone.

Funny, because, like any good parody lampooning stereotypical excesses, it contains the kernel of truth. (actually, it wouldn't really surprise me if, deep in the heart of hearts of some IT people they would aspire to that level, if they could get away with it).

At the other end of the scale are some of the unquestionably excellent, career professional, world class IT people I have attended seminars at Interop and other places with over the years. Who either are, or as some I have kept in touch with have become, the senior managers and CIO's of corporations like Citi group, Walmart, GTE and other large national and multi national companies. There can be no question the IT departments in those organizations play a vital role in the profitability of the company - if for no other reason than market forces dictate they must.

So you have to ask yourself - at what point is a company large enough to need an IT department? Where is the inflection point where it becomes more cost effective and more productive to have someone more or less full time ensuring that everyone's desktop/laptop/whatever is always up to date, virus free and has no problems accessing the shared drives and critical applications.

if I recall from the 'old days' of my early career, I used to hear figures like 'one IT person per 20 staff'. About a 5% overhead in other words. Perhaps that is right. In which case, Exetel should have three full time IT staff looking after our IT environment. Which of course, we don't. Our policy remains 'No IT department'.

Maybe it is different for ISP's, because we have Sysadmins and Network Engineers and Database Developers who, apart from their day to day jobs, can always be 'fudged' into maintaining our modest internal IT requirements as well. I think it would be either dishonest or delusional to say that doesn't happen at all in Exetel, there is always a little bit of 'creep' and overlap into that area as some tech person, trying to be helpful, somehow becomes the go-to guy for some problem.

But I can say that it happens to such a small extent that whatever extra work it creates for anyone has never been raised as a concern, nor distracted from the otherwise very full schedule our engineers, developers and administrators have. And so, whatever overlap there is, must be very small indeed.

To answer my own question above about 'when will we need an IT department?', my answer is 'never', or at least not for any foreseeable future I can envisage - because of some very specific reasons, namely;

- Initially, all our staff we engineers and everyone took their turn on the helpdesk at the start. If we couldn't make our own computers work properly, then we should not have even been there.

- But who today is not computer literate and does not have a home computer and use the Internet? Anyone who can use a home computer, install games, solve driver incompatibilities, maintain anti virus software and checks and set up all the things they want for internet has far used far more advanced skills than are every needed for the very basic, simple and straightforward requirements of using and operating a workstation in a business environment.

- anyone who can't maintain a secure, functional office desktop has that burden removed from them. How they do their work after that is up to them, but they still have to do it. No business can retain staff that can't do their work.

Sound a bit harsh? Oh well. Though it has never been a problem. If it is ever, I guess we would have to look at it and see if the policy still held.

And perhaps at least part of the reason it hasn't been, and shows no signs of ever being a problem for anyone is because who, after all, really wants the hassle of dealing with 'Nick Burns'? And by the way you're welcome!

Posted by Steve Waddington at 12:26

Wednesday, August 12, 2009

Ericsson W35 Field Trial

We have been evaluation various 'all in one' HSPA modems with the view to finding a unit that has integrated:

- wifi
- Ethernet
- ATA/VoIP

The Ericsson W35 is one of the candidates, having just about everything anyone could want in one neat unit.

We tested it from our office in North Sydney, and from my home in Perth. It is a really well made unit, the CODECs for VoIP are top notch, as I suspect are the quality of the electronics inside. So there was much interest to see how it would perform in the field. And what better place to test than our No. 1 extreme range test site in Albany.

Basically, we want a unit that will replace this:

ATA, Wifi Modem with HSPA USB Slot and USB Stick.

with something like this:

Ericsson W35 integrated unit.

The Ericsson W35 retails around \$900, but it is likely with bulk purchasing we could get the sell price down to something similar to the combined total of the individual components.

Unfortunately the excellent performance in high signal strength areas could not be duplicated with the lower signal strength at the test site.

The units signal meter registered 19% signal strength, similar to the Netcomm. But what we found was that, every few minutes, the Ericsson unit would reset itself. This could have been due to the power from the solar/battery/inverter system, although no other equipment had any problem.

The resets looked to me to be due to the low and variable signal causing the unit to think there was a problem, and reset itself to try to clear a possible fault and resync with the base station.

The unit is currently still on site for further testing with better/new coax cable and connectors, and we will also try it at another site closer to the base station and with mains power to determine just where the drop-out threshold is.

As it stands now thought, it's back to the drawing board.

Posted by Steve Waddington at 09:06

Tuesday, August 11. 2009

New Yagi Results

We have tested the new, smaller form factor Yagi antennas at medium to long distances from base stations and found there was no appreciable difference in performance. Last week I was able to test one at extreme range, replacing the first antenna we used in a field trial with one of the new ones.

New, smaller form factor Yagi installed.

7 element Yagi.

You can see the difference in size compared to the earlier model [here](#).

At 22km from town, this site is about as far as it is possible to get from a mobile phone base station and still get a usable signal. The signal strength meter on the HSPA modem, not the most reliable of measurements, reads consistently between 15 and 30%. There is no direct line of site to the base station, so I suspect there is some sort of scintillation that 'bends' the signal over the interposing hill.

Never the less, despite whatever the meter shows, the signal is constant, and drop outs seem to be no more common than in, say, North Sydney. Download speed range from a low of 800kbs to as high as 2.1Mbps, and most of the time are between 900kbs and 1.2Mbps.

The good news is that the new antenna has made no difference to the reception or performance. The performance still the same - although Bill (the property owner) thinks he might get slightly better results in stormy weather because the load factor on the mast will be less, and so it will move around less in high winds.

You can see the solar panel on the roof. The property is not connected to any utility - no power, no water, gas, land line telephone and no TV. The solar panel recharges 3 heavy duty 12 volt batteries. An inverter is used to provide 240v for some things (like laptop power), but in most cases the lights, fridge and other appliances all operate off 12 volt power. Bill was pleased to show me that the kitchen sink now has running water, recently connected to the rainwater tank outside.

I always find it interesting that while it is still quite possible (at least in Bill and Sharon's case) to live and raise a family completely 'off the grid', the one thing it is not possible to get by without is the Internet.

Posted by Steve Waddington at 09:35

Monday, August 10, 2009

Case of the resetting line card

Our astute network engineers noticed a problem a while ago. Well it looked like a problem.

The logs in one of our ESR routers showed a 'reset' condition occurring on a GE line card. The interesting thing was, despite the log entry, there was no other apparent impact to anything - no circuit drop outs, no end user disconnections, not even a packet drop on the interface as far as we could tell.

A quick trip to the CCO log interpreter tool told us "Contact Cisco TAC". Which we did. The usual sending of 'sh diag' dumps followed.

We also checked the CCO bug tool, but found nothing that would indicate the, otherwise very stable, IOS image we are using had any problems.

TAC came back some time later with the conclusion that it was a hardware fault in the line card. Funny, we had swapped that line card about six months ago for a different issue, so the one we had now was almost brand new. Oh well, these things happen.

Except that, now we noticed another problem. A second line card in the same chassis was producing the same log entry errors.

Best not to let these things drag on and then fail catastrophically, and, while we were waiting for the Cisco RMA replacement to show up, it seemed prudent to do a few things to eliminate other possible causes. Which were; replace the GBIC; move the card to another slot in the chassis; and replace one of the cards with our on-hand spare.

Personally, I was beginning to think the problem might lay somewhere else than the line cards. Never the less, the work was done, and the RMA's card arrived and was duly replaced on Friday without incident. Now monitoring over the weekend shows that the card we moved to another slot seems ok, but the new RMA'd card is registering resets in the log.

So we have three cards, one almost brand new, one absolutely brand new, apparently resetting, but when doing so not even dropping a packet, let alone the circuit. Which the TAC, who know about these things, tell us it is in fact a hardware fault.

It's sounding like a plot by a super intelligent AI to get one of us to leave the spaceship if you ask me.

What do I think it is now? Most likely, it seems to me, to be a minor IOS bug that send false reports to the log under some conditions from the line card.

But it might not be that. And it is far better to err on the side of caution than to dismiss it based on my thin supposition. One more RMA replacement it is then.

Posted by Steve Waddington at 09:26

Saturday, August 8, 2009

I bless the rains down in Albany

(apologies to Paich/J. Porcaro). Part of my trip to Albany this week was to meet with one of the Friends of the Western Ground Parrot organizers, and the project officers from Dept of Environment and Conservation about how the ARU project has been running.

The ARU trials are detailed by Chris Powell (DEC Project Officer) on the Exetel Forum here. So I won't repeat what has already been said. But what I can say is certainly talking to the people involved first hand gave me a much better appreciation of just what is involved for this sort of conservation effort.

The problem, as I understood it, with the Western Ground Parrot is that they blend in so well, and are so good at hiding, that visual sightings are almost impossible. In fact, one drawing was made in 1912, and there was another blurry and almost unidentifiable photograph from the '70's, but until 2003, there were no clear images of the birds at all.

Their actual existence was determined by their unique calls (imitated by other native birds, but distinctive once you know what to listen for). That's easy then - just go out into the bush and keep an ear out.

So you would think. Except that; they only make their calls the hour before sunrise, and in the hour after sunset. and, being in southern coastal habitats, wind and rain often obscure any calls, and, with low population densities and very few birds, they often make no calls at all for days at a time.

The traditional method of determining how many, or if any exist in an area, is for a group of half a dozen or more volunteers, organized by a DEC officer to camp out for ten days, go at an assigned area each pre-dawn and dusk, and listen for any calls the birds may make. A very labour intensive process, and even with mostly volunteers, still very expensive.

That is where the ARU trials come in. The Autonomous Recording Units contain a PEE PC, largish hard drive (300GB+), array of microphones, car battery and solar panel. They are weather hardened and designed to sit out in the bush for months at a time, recording whatever sounds there are to hear.

The units I saw were the early prototypes, very big and bulky, needing two people to carry. But even so, one field trip every few months will return a wealth of audio information for analysis.

Which is the second part of the project. Using SoundID software, known Ground Parrot calls can be used as a template, which is then matched against the captured audio, and in maybe half an hour, will return all the positive matches for the month(s) of data available.

A lot of Chris Powells' work in the latter part of the project has gone into building up the audio library and then 'tuning' the SoundID recognition parameters to return accurate matches and eliminate false positives.

But does that actually save bird one? Was the implied line of my initial questions.

No. But as I came to realize, when there are maybe just 150 birds comprising the total population, whose possible habitat extends over 50,000 square kilometers, you must have accurate data of just where the populations are to be able to know where any conservation effort is to be focused.

On the way to a Western Ground Parrot suspected population site. Mt Manypeaks rises like Olympus above the Southern Coastal Plain.

With very limited, and uncertain, funding, nothing effective can be done without knowing where to deploy the resources available.

The point of the ARU trials hence becomes apparent. In a few months several units deployed by just a couple of people can provide good scientific data, either confirming the location of a suspected population, or ruling it out, that even 50 volunteers and DEC officers could do in a year.

Low scrub heath at Waychinicup National Park. The natural habitat of the Western Ground Parrot. This location is the last confirmed sighting of a Parrot in the area in 2004.

From the left, Dr. Abby Berryman, (Research Scientist, Western Ground Parrot recovery Project) and Anne Bondin (Friends of the Western Ground Parrot), holding a portable ARU - running on 4 D cell batteries and limited storage, it needs to be determined if these units would be suitable for longer term data capture.

I will make just one other point that really hit home to me. Cats.

1080 poisoning has been effective in reducing the number of foxes. After baiting started, fox numbers declined and the native population recovered. Then when into a steep decline.

What became apparent was that the foxes had been effective in not just killing native fauna, but would also kill feral cats and their litters. When the fox numbers declined due to baiting, the cat numbers rose - and cats are far more effective predators of small marsupials and birds - and it doesn't take too much imagination to figure out just how a ground dwelling parrot is going to fare.

Cats wont take the dry baits the foxes will, and even with baits inside sausages (which seems to be the most effective to get cats to take the baits), they are still very wary. Nor are direct methods such as capture or shooting nearly as effective against cats as foxes.

Even, after much cost and effort, cats are removed from a particular area, numbers are back up withing a few years. Simply because of a) dumping of unwanted kittens, and b) uncontrolled breeding of non-neutered town cats that continually replenish the ferrel population.

It comes down to this - a choice between cats and native fauna. All moggies must be fixed.

Posted by Steve Waddington at 11:30

Saturday, August 1, 2009

Always in Threes

After the two unlikely cable cuts within a few days of each other, I made the comment that I wondered what the third thing would be.

Now I am not superstitious at all (because it is bad luck to be superstitious), but I have noticed that pretty much every time someone is foolish enough to make a comment like 'things/the network/the [some function] server has/have been running very well' the time can be then measured in hours, if not minutes, until some major problem with that particular system.

Not that anyone actually believes in such mumbo jumbo nonsense, but, never the less, best leave such optimistic statements as unsaid, or even unthought, dispense with any complacency and use the mantra 'anything can go wrong at any time'.

So, was it a surprise five days after the two fibre breakages that there would be a general power outage (caused by a, apparently, yet another, careless contractor shorting a circuit and blowing the circuit breakers on the UPS's) at one of our collocation centres in Sydney?

Only a week short of when we would be moving every other server out of there to our own, new, (and hopefully reliably UPS'd) data centre in North Sydney.

Always in threes. Well, at least that triplet is now out of the way.

Posted by Steve Waddington at 10:11

Thursday, July 30, 2009

Hobart POP

With all the excitement of a major circuit down for 36 hours, and travel to and from Sri Lanka, I forgot to mention that our new POP in Hobart is operational. So I will redress that now.

The first customer connected on the 22nd, the second on the 23rd. We have now just finished transferring the last of the Tasmanian customers still on the, very old, 'legacy' circuits connecting directly to Sydney across to the new POP.

We can now cancel the cross connects for those old circuits - which are the last remaining ATM circuits in our network (which are also I think the most expensive \$ per Mbps circuits we have). And I won't be sorry to see them go. ATM is so much less efficient than GE in terms of protocol overhead, router CPU and of course, cost.

The Hobart POP completes the Exetel National network build. We now have a presence in all Australian states and territories (not just all Mainland states and territories).

Posted by Steve Waddington at 10:00

Saturday, July 25. 2009

2nd Cable cut in 3 days

... and just what does 'premium service, 2 hours restoration' mean anyway? 24 hours, apparently.

This is the graph of one of our circuits to TW that between them carry all the ADSL1 customers connected via Telstra Wholesale:

Looks good, doesn't it.

The first outage was apparently due to rats eating through the cable.

The current outage is on the same cable, but, so far as we have been told, due to a break in the wiring cabinet. We probably won't get confirmation on that for another few hours, and the cause won't be known for sure until the incident report comes out sometime later (or I should say, we won't know the officially sanctioned and sanitized reason).

So the cable splicer guy is apparently re-splicing the broken cable even as I type. On past performance it may be another eight hours before it is fixed. I of course have no idea how Telstra field techs do these things, but it seems our particular fibre in the bundle is always the last one to be done - because the restoration notice (for the whole fibre bundle) arrives half an hour after we see the service restored, from which I conclude ours must be one, if not the last one, done.

Two major day long breaks on a '5 nines' service in three days. How very, very unlucky. Maybe Telstra is looking to take Optus's title of 'World's unluckiest Carrier'?

I wonder what the third thing will be?

(and so much for my 'karma positive' blog idea)

Posted by Steve Waddington at 18:07

Thursday, July 23. 2009

Exetel Scholarship in Unified Communications

Over a month since my last blog article. That's terrible, what have I been doing? Well, I have started to write a blog article many times, but in each case it just seems to descend into anger and frustration over the sheer unreasonableness people, suppliers or pointless, parasitic, self serving government bureaucracy.

See?

I decided to try and make blog posts 'karma positive', and so, twenty odd drafts have been consigned to a publish date of 2222, or until I can figure out a way to rewrite them in some positive light.

Anyway, enough about that.

I am in Sri Lanka this week for an in-house training course that I conducted over the weekend, and to present the two inaugural Scholarship grants to the Sri Lanka Institute of Information Technology.

On Monday Rohan and I, along with Dr. Kapurubandara (Dean/Student Affairs) completed the selection process started by the SLIIT by interviewing the final six candidates. These were the best performing students based on their first semester results and their previous high school academic record. Being all distinction and high credit students, there wasn't much between the candidates in terms of grades, and the eventual, not easy decision, was based on the two candidates to whom we thought the scholarship money would make to most difference to their home and personal life.

We are providing two fully funded scholarships per year, for the duration of the four year degree program. Which means after four years, there will be eight scholarships grants per year.

The SLIIT campus is in Malabe, a suburb about 20km to the east of Colombo. Situated at the top of a hill, the ten story main building offers spectacular views across the valley to foothills and then three levels of rising mountains that are the central mountain range of Sri Lanka.

So, we decided on the two, and a small award ceremony would follow on Wednesday at 10am.

We arrived a little early on Wednesday, as the same time as the SLIIT President and Chairman, Professor Karunarante. As we walked across the courtyard to the main entrance, it really hit home to me just how important and seriously Sri Lankan's view tertiary education. Every student, there were maybe two to three hundred in the courtyard at the time, stood up and stopped talking as Prof. Karunarante (and us) walked across the courtyard and into the main building.

"They are standing to show their respect" Rohan informed me "University education is very important in Sri Lanka"

It was an amazing thing to see. I just couldn't imagine anything like it in Australia. We clearly take things like a University education too much for granted, and don't appreciate what it might mean to someone whose whole extended family might have had to scrape together the equivalent of a years wages to pay their tuition fees.

The 'small' ceremony was attended by about 50 students, members of faculty, the families of the scholarship winners, and journalists from four radio stations, two TV stations and the two national news papers. The podium bristled with microphones like a UN press conference. No pressure for the short speech I was going to give then.

An excellent job by the SLIIT PR department, to be sure.

The scholarships were awarded, the cheques were handed over. There were two very happy students, and we have the start of a great ongoing relationship with the SLIIT.

Posted by Steve Waddington at 13:19

Tuesday, June 16. 2009

The Next Big Move

The first of three new 1Gbps Internet bandwidth links is now up. All three were physically connected on Friday and now we are waiting for the end and third circuit to be activated with a live IP address.

(or as I joked to Jing, for the carrier to take the clamps off the fibre to let the photons through)

With one circuit up, we are now testing the BGP setup and connectivity to the Internet with the staff IP address blocks. Then when we are happy with that, we will start announcing some server addresses and then progressively add in end user blocks over the week.

The end result is that we will replace 1.9Gbps of bandwidth in NSW with 2.4Gbps.

While such a gross move of bandwidth is not something to be done without a lot of anxiety (which no amount of preparation and contingency planning can ever completely eliminate, and probably quite rightly), there is never the less some cause to feel slightly more mellow in this case.

First, the new circuits are on a, currently, completely unused router (our backup 6500). So we have the comparative luxury of no chance of affecting any production traffic with whatever mucking around that is inevitably needed to get the circuits, IP and bgp up and running.

Second, that 6500 used to be the production border router, so it is still cabled in with known good circuits to the production LAN backbone. That means we don't need to change the topology at all, or make even any minor changes to the network as it sits.

Third, once we are happy with BGP, we can simultaneously announce address blocks on both circuit, prepending the old circuit and load test the new one with no disruption to any production traffic. And if something unexpected does happen in that regard, the fix will be only a matter of seconds to change the circuit preference.

Forth, we can if absolutely necessary, extend the cut off time for the old circuits past the target date of the 20th. It would only be as a last resort if something went badly wrong (and there are no indications of that at this stage). But even with out that, with the circuits active today, we are right in the middle of the critical path schedule.

Finally, the network team have done a comprehensive job with planning, and with my 'mother hen' plan-for-the-best-but-expect-the-worst hovering, it is hard to see what possible problems we have not allowed for.

Which maybe is what bothers me. Where is Saffron when you need her.

Posted by Steve Waddington at 12:24

Friday, May 29. 2009

Twitter

As an experiment I have been using twitter to keep interested parties up to date on an event of some possible impact.

I am struggling though to see what the claimed press hype on the Twitter home page is all about. Of course I can't comment on it's use as a 'social network' tool - which is something I have only a very vague idea of.

Twitter seems to be most useful (commercially) a way to get a short message to many interested parties, at no cost. However, I find the 140 character per message limit a little, well, limiting. I think a well managed forum probably beats it there.

Also, no pictures. I wanted to put some before and after shots of the new computer room as work progressed, which I could certainly have done in a blog.

It seems to me that if you cut Facebook back to the bone, then that is really what Twitter is. So why not just use Facebook, which does everything Twitter does, and has far more features?

Still, Wired, Newsweek the New York Times, and other journalists that wax so positive about it can't all be wrong, can they? I must be missing the point somewhere.

Posted by Steve Waddington at 11:50

Thursday, May 28. 2009

The Big Move

We are in the process of moving offices from the place we have occupied for the last five years to the new office in Walker Street.

I am going to test out 'twitter' as a way of keeping everyone up to date.

Anyone interested can follow our move blow by blow here:

<https://twitter.com/SteveWaddington>

Posted by Steve Waddington at 17:07

Monday, May 25, 2009

UC, er I mean VoIP, or UC... whatever

We all must have Unified Communications and our workforce must be Mobile Enabled, and so went a new bevy of buzz words at Interop. No one is mentioning 'three play' or 'quad play' (never called four play for some reason) anymore, so I suppose that is a plus.

Having had a think about what I say last week, and what we currently have, there is actually nothing we are missing from our product line up that doesn't meet or exceed the 'best of the best' solutions on offer at Interop. Well, actually there is. One big thing - A nice and slick user interface.

What we need is a web based client that will integrate all VoIP services to the PC or mobile device.

While all the controls are there now in the user faculties section, I think we just need to bring them together into a smoother 'control panel' type interface.

The best one I saw at Interop was by Shoretel. Who had another 'interesting' thing to say. Towing the super safe 'don't sue us' line that all American VoIP manufacturers seem to have adopted, they also were adamant that VoIP services should/must not be used where:

- latency is greater than 150ms
- jitter is greater than 15ms
- packet loss is greater than 0.5%

Well, as you can see from an earlier post, my disappointing experience at the LAX Qantas flight lounge was:

Ping statistics for 220.233.0.1:

Packets: Sent = 47, Received = 34, Lost = 13 (27% loss),

Approximate round trip times in milli-seconds:

Minimum = 289ms, Maximum = 693ms, Average = 450ms

Jitter in the 200-300ms range, almost 1/3 of packets lost and latency 4-5 times the 'must not exceed' value.

I called home in my last test of wifi VoIP for the trip on that connection. The quality wasn't great I have to say, lag was very noticeable and the tone was a tinny. But - the call completed, I spoke with my wife for about 10 minutes and the kids (who are far less tolerant of low quality voice) for a minute or so each. There was no noticeable echo and the speech on both sides remained legible throughout.

So there you go. VoIP is usable in the real world, on links that are about as bad as they can be. No matter what the American manufactures want to tell you.

Posted by Steve Waddington at 17:06

Sunday, May 24, 2009

Interop Best of Show

Just a few quick highlights from Interop while I sort my notes and get the many ideas the show sparked into some sort of order.

What I can say though, is that I can see a clear path from here to how we can make the Exetel network substantially more extensible.

More of that later.

The best booth was by far the Xirrus display with their boxing ring and live boxing show, where prize fighters duked it out to determine the advantages of matrixed wifi over wired LAN's. Great technology too, by the look of it. Xirrus supplied wifi for the whole conference with 18 nodes covering the whole 5.5 million square feet.

Passions run high as Melissa Littejohn and Ashlei Fesolai battle it out over wifi vs wired LAN's.

Cloud computing featured heavily. 'So just what is cloud computing?' I asked one of the reps at a vendor stand. He rolled his eyes and suppressed a groan, either at the stupidity of my question, or my non USA accent, or both, but answered it anyway. So at least now I know what cloud computing is.

The biggest technology, championed by every major, and most minor, vendors however, was 'virtualization'. Now that I know something about, having followed VMWare with interest from their very early days, and trialled it on servers as early as circa 1999-2000.

Looking at how hard IBM, Dell, Cisco and every other hardware and OS vendor are flogging it, it is a technology if not come of age, then about to very, very soon.

Also of interest is that everyone (US carriers I mean), have dropped the MPLS 'Saviour of your network' crap (who needs to know, and who cares), but now, suddenly, their WAN services have become (or always been, according to the marketing) 'Virtualization Optimized' or 'Virtualization Improving'. And how would they do that exactly? Well, the details, unsurprisingly were scarce, since apart from having a network with no contention, there isn't a lot else the WAN provider can do. But it does go to show just how importantly CIO/CTO's view it, that the carriers feel compelled to use it in their marketing at every opportunity.

What a pity that so many Australian WAN suppliers have all those glossy expensive MPLS brochures still to get rid of, and their reps have just learned to spell it (but probably not what it actually means - will have to ask one next time I hear it). I guess we will still be hearing the tired old BS claims about it from their reps for a while yet, before they catch on to the international trend.

But as far as virtualization goes - it is certainly something we will be looking into very closely in the next few months. Especially as some of our older servers come up for 'retirement'.

Posted by Steve Waddington at 11:47

Friday, May 22, 2009

Leaving Las Vegas...

... Shame and loathing in LA.

I was fortunate enough to get an FF upgrade on my flight back to Sydney to first class. Not bad I thought, since I had a four hour wait between my flight from Las Vegas till the 11pm flight from LA. The wait should be a little easier in the salubrious surrounds of the international first class lounge.

Yeah right.

Maybe the wait for the transfer bus to the Tom Bradley International Terminal, running every half hour, didn't put me in the best of moods.

"Have a walk about while you wait, mate" said the young Aussie Qantas attendant at the transfer desk. "There's bugger all to do once you get over there"

He didn't actually say 'bugger all' but I got the idea he wanted to.

I have traipsed acreages of hallways in the last few days, and seen everything I ever wanted in the way of airport terminal shops already. So I just took one of the plastic and vinyl seats, the near pinnacle of sadistic seating design, to wait the 25 minutes for the next bus.

The transfer bus slowed down long enough for most of us to get off somewhere near the back of the terminal, near the waste disposal bay, and most of those of us that made it off managed to find our way to the main terminal gates with no one to direct us and no obvious signage. Maybe that was due to some deconstruction work that seemed to be going on, tastefully masked behind tattered plastic sheets and duct tape.

After solving the puzzle of finding the way to the flight lounge - one small sign on an elevator hidden around the corner of the security staff tea room - I was greeted by a very pleasant attendant who checked my boarding pass and ushered me into the lounge.

Ah, a sigh of relief, I can relax in comfort for a while now, I thought.

My expectation didn't, and still doesn't, seem unreasonable to me. The Emirates lounge at Perth airport is first rate, so too is the NZ Airlines lounge in Auckland and the old Ansett lounges when they were still flying. The Singapore Airlines lounge is just fantastic, and the Thai airlines lounge is not so shabby either.

And all those are the business class lounges. Nor is there anything to complain about at the shared facilities lounge in Sri Lanka.

So that was what my expectation was set by. Qantas and OneWorld, first tier airlines. LAX, a major international airport. First class lounge - going to be nice.

No.

It is crap.

One small step up from the abysmal dives American domestic carriers call business and first lounges - where you are wedged shoulder to shoulder to use your one complimentary beverage voucher (and still have to tip anyway).

But still crowded. The drinks are free, and at least getting them yourself means you are saved the dilemma of tipping. Though they are nothing special - \$20 a bottle Wolf Blass, some Nappa Valley Merlot that tastes like cordial. Usual mix of spirits and some run of the mill international beers.

Nor is the food much better. A very small 'help yourself' area with a few overcooked and dried out selections. Some cheese and biscuits, still fresh from their plastic wrappings and a 'make it your bloody self' espresso machine.

Ok, well those LV buffet's take a bit of getting over, and I wasn't that hungry anyway. I'll just use the restroom. Ugh, maybe not. The lounge has a capacity of maybe 100 and there were about 70 people there. The mens room had two urinals and one stall. I mean COME ON! The one thing I can say that is positive about American airports is that, most of the time, the toilets are clean. Not so in the first class lounge though it seems. Maybe we filthy foreigners don't rate having clean restrooms (or maybe they are that way to show the Americans traveling on a non-domestic carrier how much better off they are with their own).

I'll just sit and catch up with some email then. The first trick is to find a seat where I can plug my laptop in - no mean feat, since there are no work areas as such (maybe frequent first class travelers don't lower themselves to that sort of thing). After a bit of a wait, I snagged a seat near a pillar with a power point that someone had just left.

As I sat down, one of the (two) lounge attendants gave the micro table in front of me a cursory wipe with a cloth, helpfully smearing the remains of whatever the last three people had left into a flatter surface I could at least now balance my laptop on.

Wow, wifi is actually free! For a moment there I thought I would have to buy access. Hurrah, I can now download my email.

Connection to server lost.

pop3.exetel.com.au host not found

```
C:\Windows\system32>ping -n 50 220.233.0.1
```

Pinging 220.233.0.1 with 32 bytes of data:

Request timed out.

Request timed out.

Reply from 220.233.0.1: bytes=32 time=365ms TTL=52

Reply from 220.233.0.1: bytes=32 time=357ms TTL=52

Reply from 220.233.0.1: bytes=32 time=446ms TTL=52

Reply from 220.233.0.1: bytes=32 time=511ms TTL=52

Reply from 220.233.0.1: bytes=32 time=389ms TTL=52

Reply from 220.233.0.1: bytes=32 time=416ms TTL=52

Reply from 220.233.0.1: bytes=32 time=531ms TTL=52

Reply from 220.233.0.1: bytes=32 time=461ms TTL=52

Request timed out.

Request timed out.

Reply from 220.233.0.1: bytes=32 time=370ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=364ms TTL=52

Reply from 220.233.0.1: bytes=32 time=571ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=459ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=365ms TTL=52

Reply from 220.233.0.1: bytes=32 time=554ms TTL=52

Reply from 220.233.0.1: bytes=32 time=503ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=693ms TTL=52

Reply from 220.233.0.1: bytes=32 time=514ms TTL=52

Reply from 220.233.0.1: bytes=32 time=452ms TTL=52

Reply from 220.233.0.1: bytes=32 time=537ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=490ms TTL=52

Reply from 220.233.0.1: bytes=32 time=403ms TTL=52

Reply from 220.233.0.1: bytes=32 time=289ms TTL=52

Reply from 220.233.0.1: bytes=32 time=490ms TTL=52

Reply from 220.233.0.1: bytes=32 time=355ms TTL=52

Request timed out.

Reply from 220.233.0.1: bytes=32 time=476ms TTL=52

Reply from 220.233.0.1: bytes=32 time=369ms TTL=52
Reply from 220.233.0.1: bytes=32 time=526ms TTL=52
Reply from 220.233.0.1: bytes=32 time=524ms TTL=52
Reply from 220.233.0.1: bytes=32 time=423ms TTL=52
Reply from 220.233.0.1: bytes=32 time=448ms TTL=52
Request timed out.
Request timed out.
Request timed out.
Reply from 220.233.0.1: bytes=32 time=337ms TTL=52
Reply from 220.233.0.1: bytes=32 time=338ms TTL=52
Reply from 220.233.0.1: bytes=32 time=570ms TTL=52
Reply from 220.233.0.1: bytes=32 time=412ms TTL=52

Ping statistics for 220.233.0.1:

Packets: Sent = 47, Received = 34, Lost = 13 (27% loss),
Approximate round trip times in milli-seconds:
Minimum = 289ms, Maximum = 693ms, Average = 450ms

Qantas, you are a class act. In a country where bandwidth retails for \$2 per Mbps, you use a dial up connection for your flight lounge. Good on ya, mate.

Are Qantas tickets cheap? Sure, mine was an FF upgrade. But you look at business and first class ticket rates for any airline between Australia and the US, and the Qantas prices are either the highest or nearly there. Usually always 30-40% more than Singapore, Malaysia or one of the Japanese carriers.

I feel embarrassed to be an Australian if this is the best we can do, and I feel bad for my country that our national carrier cheapskates premium travel so much. Why would anyone, paying for a premium ticket class, not choose Singapore Airlines, JAL, or even travel the other way via Emirates or one of those other superb airlines?

Oh, and now the lounge is full and I must share my micro table and one power outlet with a fellow suffering traveler. Camaraderie in the trenches an all that.

What a disappointment.

Posted by Steve Waddington at 13:44

Monday, May 18, 2009

Travel to the US and VoIP

On my way to Interop in Las Vegas I have had the chance to try out our new VoIP client for the Nokia N96, and how our new Asterisk based PABX would handle the latency.

First stop off was Auckland Airport. The bad thing about the free airport wifi is it still requires registration via a captive portal before you can get external Internet access. That is just fine for a VoIP client on a laptop, but a pain on an N96, since the wifi network has to be connected first before the VoIP app is activated. Took about six goes, but I got there eventually.

Called home in Perth, no problems, pretty much a MoS 3.6 call. Dropped out after about 10 minutes though, but no big deal, my plane was boarding anyway.

Next stop was LAX, and the wifi at the LAX Hilton. The phone found the wifi network and the VoIP client registered straight away. I started dialing with +61, but then remembered I was getting Sydney 'dial tone', so just needed my 089xxxx number.

No problems with that call at all either, slightly better than the NZ one, maybe less contention on the local network. I gave it a MoS of 3.8. The call also dropped out after 10 minutes, but that was because the hotel only provides 10 minutes free, before it diverts you to an active portal to register and pay.

Then it was on to Las Vegas, and the free wifi network at Interop. The latency of the Interop LAN, just to its WAN boarder was 110ms, and another 60ms to LA - due to every vendor and his dog flooding the network to demonstrate their high speed WAN applications, and other deliberately flooding it to demonstrate the effectiveness of their QoS (look, see how crap the network is, and with a flick of the virtual switch, now it is better). Unfortunately my phone wasn't registered with any QoS network vendor, so I just had to accept the latency, packet loss and jitter hit.

It was just on 8am Perth time when I called, and the voice quality was just fine, but there was a lag of maybe a full second. Again the call dropped out after 10 minutes - maybe someone is trying to tell me something - but when I reconnected the lag was much less, in fact not even noticeable. So I suspect some vendors intelligent traffic management system noticed my VoIP packets struggling, got hold of the data stream and did something with it that caused the drop out and then significant improvement.

Anyway, the first call I rate at MoS 3.2 and the second MoS 4.

Now, on to the conference.....

Posted by Steve Waddington at 17:06

Friday, May 15, 2009

Sales Engineer Training Course

I was determined that the next training course I ran would be 'virtual'. I must be four years ago now since I last ran a course for engineers, and in that one I used email lists and shared network drives to disseminate all the information to the attendees.

It was an experiment for me to see just what would be possible in terms of conducting a structured course completely on line with attendees able to be anywhere in the world.

There is no doubt in my mind that that can be done. But. Time has slipped by, and I am taking the 'lazy' way out and getting our engineers and some sales people into the office tomorrow for a 3 hour course.

This is the agenda:

9:30 - 9:35 Introduction

- course objectives - 2 deliverables and one practical exercise
- scoring

9:35 - 9:45 Sales Cycle

9:45 - 9:55 Engineering Cycle

9:55 - 10:25 workshop: design a sales/engineering process

10:25 - 10:30 Team 1 presentation

10:30 - 10:35 Team 2 presentation

10:35 - 10:36 team/individual scoring

10:36 - 10:45 Experience with bad engineering habits

10:45 - 10:50 Experience with bad sales habits

10:50 - 11:00 Productive dialog and meetings will solve most if not all problems

11:00 - 11:15 workshop: teams conduct a sales engineering meeting for a hypothetical sale

11:15 - 11:20 team 2 minutes presentation

11:20 - 11:25 team 1 minutes presentation

11:25 - 11:26 team/individual scoring

11:26 - 11:45 Providing and exchanging clear and concise information

11:45 - 11:55 example sales calls to an interested buyer, per team

11:55 - 12:19 workshop: each team to produce costed solutions

12:19 - 12:24 team 1 presentation

12:24 - 12:29 team 2 presentation

12:29 - 12:30 team/individual scoring

12:30 END - Lunch

A packed schedule to be sure - but when I am the main presenter, it needs to be busy to overcome my otherwise boring personality. Dammit Jim, I'm an engineer, not an entertainer.

Anyway, all our people are bright, so I sure they will handle it.

I am expecting it to be a productive use of a morning. There are always three results that should be aimed for in any course with your own company:

1. the obvious one - effectively importing knowledge to individuals
2. the presenter finding out what the individuals know/can learn
3. each person gaining better understanding of their colleagues

Which I find that taking what any sensible trainer would consider a two day course and compressing it by a factor of 4 will best achieve. Because my experience has always been that the most is learned when the pressure is greatest and the time is shortest.

Posted by Steve Waddington at 14:19

Thursday, May 14, 2009

A very convenient Scandal

Is there anything a government likes more than some sort of scandal to distract everyone from contravention policies and budgets? Sex or sport scandals are best, and combining both is a double bonanza.

I have to say I have been a long time watcher of 4 Corners, and I can't recall ever seeing a more sensational or less balanced piece.

The timing couldn't have been arranged better if the Prime Ministers own media department had coordinated it.

Posted by Steve Waddington at 15:18

Wednesday, May 13, 2009

Filter Trial Sales Disaster...

.. for Exetel ill wishers and competitor stooges.

I waited two weeks to make sure I could get comparable figures from last month, and that any 'fall out' would be fully apparent.

And the results are (drum roll) - an eight percent increase in ADSLx sales since we ran the trial.

It is very unlikely to be a result caused by the trial itself. Most probably it is a combination of normal and seasonal variance.

I am so sorry - not - to shatter the ego's of the rabid mob (four is a mob, right?) baying for Exetel to collapse because of our filter trial.

Oh well. In the words of Lennon and McCartney "life goes on, brah!"

Posted by Steve Waddington at 17:23

Friday, May 8, 2009

A Very Convenient Recession

Just read this: <http://www.abc.net.au/am/content/2008/s2564238.htm>

So what is the score now? Thanks to the stimulus packages, no Australian child now needs to live without a plasma TV and an XBOX. Some retail staff have a reprieve on their redundancy, and a medium sized warehouse in China has cleared some of its stockpile.

Well, that is certainly a better way to justify a huge budget deficit than spending on needless luxuries like health care.

How convenient though that the GFC can be used as an excuse for broken electoral promises, just as it can for large corporations to justify work force reductions without too much union protest.

It doesn't seem to be in anyone's (any one that matters that is) interest for a recovery just yet.

Posted by Steve Waddington at 11:48

Monday, May 4, 2009

Filters and Blacklisting

With all the vehement who-har over our recent filter trial last week you would think those same people would go apoplectic with protest over an even more uncontrolled, unaudited and secretive form of censorship that is far more widespread. Yet have you heard any voice raised, ever, about the many private organization blacklists that infest the Internet?

No. Well, not from the crazies to get so mad about a child pornography filter list anyway.

In fact, I wonder if it is any coincidence that Exetel DNS servers now appear on two of these blacklists? - I doubt it.

Perhaps the cries of 'We hate censorship and we hate you' really mean 'We hate not being the ones who control censorship and we hate you enough to censor you which we will do using our own clandestine methods'.

Blacklists have been around since at least the mid 90's, originally as an attempt by the Internet 'community' to control poorly configured and compromised servers. Later being used most often for spam. They work by publishing a list of server IP addresses that sysadmin can get (or set up an automated process to get), that they load onto their own server to prevent email or other traffic being received from or sent to the blacklisted server.

The idea is (or was) that if all sysadmins use the blacklist, then the 'bad' servers would find themselves isolated from the rest of the Internet, and the sysadmin controlling them would have to take some action to fix them up.

It was always up to each individual sysadmin to decide to use blacklists or not, and which ones, because there are many. Most claim to be reputable and to operate ethically (however that can be defined). And it is arguable that they may indeed have in theory some beneficial effect.

In practice however, I can tell you, they cause far more problems than they resolve. Consider the answers to these two questions:

* Who sets up the blacklists? There is no central authority or governance. They are created by individuals or groups based on their own desires and beliefs

* Who decides what is blacklisted and why? Often the list creators publish the criteria for blacklist entry. But who can say if that is followed or how entries are actually made. Methods range from anyone being able to nominate an IP address to supposed 'honey pot' captures of spam. Actual checking is time consuming, and since often blacklists are provided 'free' it is unlikely sufficient checking is done in all cases, and possibly no checking is ever done.

It used to be that all such blacklists were free, and run by volunteers. I think that was the point that blacklists had the most credibility - at least there was no monetary gain for the maintainers (putting aside personal grudges they may have had that might cause an entry to appear on a list).

Although now that I think about it, I do recall having a heated phone call with one such 'do gooder' operator of a list a few years ago, who threatened that he would 'black hole you (the ISP I was working for) into oblivion' if we did not comply with his wishes to block a particular customer server he believed was the cause of some spam. I kid you not, he was insisting that I illegally block a paying customer on his say so, and threatening to severely impact the operation of the company I worked for and jeopardize the livelihood of the 20 or so people who worked there.

'How can you say you will do that?' I asked him

'Because I hate spam and I will do anything to stop it' was his shouted reply

The sad thing was, he actually had some power to do that. Not as much as he obviously liked to think, but enough to make a reasonable impact on the business so that his threat could not be taken lightly. Who gave him that power? Who regulated it? Who ensured it was not abused? No one, he was a self appointed vigilantly who relied on the fact that many sysadmins were ignorant enough and slack enough to think that his blacklist was a good idea.

So there you go. Blacklists are ineffective at best, and cyber terrorism at worst.

The latest block lists are now run on a much more commercial basis. The unfortunate that appears on the list must pay an 'administration, not for profit' fee to be removed, and anyone wanting to find out the reason for the block must also pay some sort of fee.

We have moved from cyber terrorism to extortion.

But where are all you freedom of the Internet proponents now? Where have you been for the last 15 years while this travesty against free speech has been carried on in its many iterations?

I would make a bet that in many cases the same self appointed guardians of Internet freedom are the maintainers of, or at the very least contributors to, the various blacklists.

Anyway, enough about those wankers.

How do the sysadmins get away with using blacklists in the first place?

If a sysadmin or IT manager went to their MD and said 'I am going to put in place a list that will block access to parts of the Internet. I don't know who compiled the list, I don't know what criteria they use, but it will prevent some legitimate emails getting through, and black hole access to other legitimate sites that our staff might need to access. Why am I doing it? Well, there MAY be spam coming from the an address on the list, but I can't be sure and have no way of making sure'. How long do you think they would stay in that job?

I can only imagine that what actually happens is something like, the sysadmin puts the list in place without telling anyone because (being lazy and ignorant) thinks it is a good idea. Email or sites are blocked, the users or management ask why, and get told 'That site is blacklisted for spam (or whatever first plausible reason that pops into the sysadmins mind).

No further questions are asked, because, of course who wants to waste time with a half hour explanation of technical gobbledegook that no one understands. Or perhaps, who wants to look ignorant by asking why? He/She is a sysadmin, they must know what their doing, right?

Well here's a tip; If your sysadmin tells you that some email or site is blocked because of an external blacklist they have unilaterally decided to subscribe to (or failed to remove), sack them. Any doubts you had about their competence are well founded, and the unease you have always felt about them being responsible for a vital part of your business is vindicated. Sack them now before they do any more damage, then go home early because you have probably made the most significant contribution to your profitability for the year.

Posted by Steve Waddington at 09:37

Thursday, April 30, 2009

Filter trial causes traffic slow down...

... in the Sydney Harbour Tunnel!

Which is a humorous comparison some people have made on the Exetel forum to some of the 'false negative' reports received about the filter trial.

An analysis and blind test results are here: <http://forum.exetel.com.au/viewtopic.php?f=4&t=31857&p=243700#p243700>

To coin a phrase, 'The Harbour Tunnel Effect' could be used whenever one event is held responsible for another event that happens at the same time, but is completely unrelated to it.

Posted by Steve Waddington at 10:11

Wednesday, April 29, 2009

Why run a filter trial...

... that isn't even the preferred filter method?

I'm glad you asked.

I will preface my explanation by saying that 'we' all know no content filter is ever going to be able to filter anything someone really wants to access. While any filter system is a crazy idea, if we HAVE to have one, would it be preferable to have one that is a) far less effective at stopping anything, and far more complicated and expensive, or b) one that is as effective as possible, and comparatively simple and inexpensive?

Hmmm, let me think about that. For me, I would go for 'b'. But in case it isn't clear, let me put it in end user terms; would you rather a) pay \$2 per month extra for your Internet in exchange for the minor inconvenience of perhaps not being able to access some sites without a work-around, or b) pay 50 cents a year extra for your Internet in exchange for the minor inconvenience of perhaps not being able to access some sites without a work-around?

(NONE AT ALL. EXETEL IS THE GREAT SATAN FOR EVEN MENTIONING IT. RAGE SPIT FOAM. Yes, thank you all you crazies, as always, you have served your purpose.)

My guess is you would say 'b' too.

Why then is an opt in/out filter system, option 'a', the 'preferred' solution? It can only be because of the very negative public opinion the work 'mandatory' evokes. 'Optional' on the other hand, to someone who doesn't know otherwise, sound much more palatable.

I wont re-hash my argument of why optional filtering makes no sense (or rather, even less sense). Lets just take it as a given that someone with no other knowledge or information is going to prefer something bad that is 'optional' to something bad that is 'mandatory'.

So much so in fact, that I believe that if content filtering were to be made non-optional then the labour government, ever concerned over public opinion (as demonstrated by the plastic bag back down), would have to scrap the project.

When you are well briefed on the technology, the implications, options and costs (as I like to think I am, but no more than anyone with a technical background), and work though the political arguments presented, it becomes clear that there can be no POSSIBLE reason for opt in filtering being preferred other than it is the only way the scheme could possibly get through.

There are two 'battles' to be fought. The obvious one is whether the content filter itself. But disagree as much as you want, the fact is slightly more than half of you voted for this government who therefore has a mandate to fulfill its promise.

The less obvious one, and the one everyone has been successful distracted from, is the type of filtering system, which brings us back to option a) or b).

With official trials focusing on opt in/out filter systems, what evidence is there that anything else is better/cheaper? What debate is there about alternatives? And what can make the 'general public' (who with no other input will think 'optional' is better than 'mandatory') aware of the difference?

I guess someone should run a trial that will do that.

Posted by Steve Waddington at 11:36

Tuesday, April 28. 2009

Content Filtering - Opt In/Opt Out

It doesn't make sense. For the sake of argument, just put aside the difficulties of allowing every end user to opt in or out of a content filter and consider what possible reason there could be to do such a thing?

If the government is serious about content filtering, for the strongly alluded purpose of preventing child porn and other illegal content, then how can there be an option that allows people not to use the filter? Its like legislating against drunk driving but allowing the driver an option to stop for an RBT. Or, 'I'm sorry Mrs Jones, we would like to take your husbands murderer to jail, but he has opted out of a trial'.

There is no doubt murderers, drunk drivers and other criminals would opt not to be caught if they could. So who do you think would choose to opt out of a filter system that blocked illegal content?

If a government imposes a content filter system, then is also must have the backbone to make it mandatory, otherwise it is nothing more than senseless grandstanding of weak willed leadership to make it look like they are doing something to fulfill an empty election promise (and wasting lots of money in the process).

Oh... damn.

I guess we will end up with an opt in filter system then.

Posted by Steve Waddington at 09:31

Friday, April 24, 2009

Content Filtering

We have just finished testing out, prior to a live trial, a content filtering system put forward by Watchdog from New Zealand.

Yes, yes, I know content filtering is a travesty of our rights (...to watch porn), thought control by fascist regimes, a return to the dark ages, etc. And, as I have argued elsewhere, not something I agree with - it being inappropriate, in my opinion, to use technology as both scapegoat and cure for a social problem.

Never the less, should sibling Rudd mandate it, we, the people who will have to put it in place at Exetel, need to know how to do that. So, agree with it in principle or not, a trial of the technology before that happens is a sensible thing to do.

The Watchdog folks first contacted me maybe eight months ago, and after three or four terse brush off replies to their emails, their polite perseverance lead to an eventual meeting at our Sydney office. Mainly because I figured I probably did need to understand what could be done with regard to filtering, and the solution on their web site did look workable.

"We block content using BGP announcements to redirect traffic to a block page" They told me.

"Ah ha!" I pounced, "how is that any different to a 'null0 route' and how do you avoid overblocking?"

The went on to explain exactly how that was done, and I have to admit, it is a pretty elegant solution to the problem.

Peter, please comment or email me if this is wrong; The way it works is a BGP session is set up between the ISP and the filtering server. The filtering server has a list of sites (supplied by the government or whoever), and resolves them to their IP address. It then announces those addresses as /32 prefixes to the ISP router.

As you can see, the ISP traffic requesting access to any site on the list will go to the filter server. The filter server acts as a simple proxy for any URL that is NOT the blocked site specifically, so the end user sees no difference. But for the URL that matches the block list, the server instead returns a page that informs the user the site contains illegal content and has been blocked.

Very simple and effective. What I like most is it imposes virtually no extra load on any part of the ISP network. At the same time it neatly side-steps overblocking in a very practical way, and it is hard to see how even a low end server would not be able to meet the demands of a very large ISP.

It places so little load, and needs so little bandwidth, that for the trial we have created a tunnel between our border router and a Watchdog filter server in New Zealand to make them look adjacent. The Watchdog filter server is using a filter list from a similar service in the UK, as defined by the government over there. Which, probably, will be close to the filter list the Australian government will use to protect us naive and innocent Australians.

Currently the filter server is offering us 198 /32 prefixes and testing shows the redirection is working just fine. When we announce the live trial, it will be interesting to see what feedback we get, and who and what it really affects - if anyone or anything.

Update: I have just started the forum thread for the trial here: <http://forum.exetel.com.au/viewtopic.php?f=4&t=31829>

Posted by Steve Waddington at 10:48

Wednesday, April 22, 2009

NeoNaziBay

Ok, that was a cheap shot. But the attitude that PirateBay seem to share with all pro-copyright infringer's is worthy the level of denial (we can not lose!) that led to der Untergang.

So the PirateBay 4 don't disagree that unauthorized distribution of copyright material is illegal, nor that people are wrong to do it. Rather, their argument is that by facilitating that and providing in their words '... the worlds largest bittorrent tracker' they are doing nothing wrong. Oh, and the fact that they sell advertising and make (at least enough to fund an expensive legal defense) profit is just bye the bye.

"I didn't commit the armed robbery your Honor, I just provided the plans, loaded the guns, watched it happen and then drove the getaway car AND [the crux of the argument] one person did withdraw money from their own account at the same time, it was really them I was there for. Sure I got _paid_, but that was just for services rendered, it had nothing to do with the robbery itself"

How strange that the Swedish court had a different view.

The denial continued with this quote from their blog:

"The site will live on! We are more determined than ever that what we do is right. Millions of users are a good proof of that."

Millions of good users? Really? Who would they be? Would that be the millions of users who download Linux ISO's or other legally sharable content? Let's see... hmmm, there are 43 seeders for Mandriva linux, which is the most popular torrent that is clearly non-copyright infringing. There are 328 seeders however for "Mac OS X Leopard 10.5.6 Full Retail", which I somehow think is copyright infringing.

Or are they the millions who download other things, like:

Heroes.S03E24.I.Am.Sylar.HDTV.XviD-FQM.avi - 43,210 seeders
Fast and Furious 4 (2009) TELESYNC XviD OPTiC-MFDâ,,ç.avi - 21,843 seeders
Yes.Man.2008.DvDRip-FxM - 8,429 seeders
etc, etc..

So the 'millions' seem to be the people who actually committed the robbery - AKA thieves - who are saying that _they_ think it's alright, and on that basis PirateBay - AKA the aiders and abettors - can determine that they are in the right. Well that's alright then.

Regardless of the outcome of the appeal I just don't see how any argument based on solely on denying responsibility for something they are clearly responsible for and such cyclic illogic can stand up in the long term.

Now the groundswell of support for Piratebay is from the "Millions" who are the pirates (funny that), but who are only the very vocal minority. One day, the house of cards of P2P copyright infringement is going to come crashing down because, at the basic level most people are a) sensible and b) honest, and the case against Piratebay and the appeal, successful or not, just helps to raise the issue to the level were more sensible and honest people will notice.

Or maybe that loud rumbling really is the support of millions of honest users, and not the Russian tanks closing in on the FÃ¼hrerbunker after all.

Posted by Steve Waddington at 10:54

Tuesday, April 21. 2009

Lower costs, but more of them

From John's blog today:

http://www.nytimes.com/2009/04/20/business/20isp.html?_r=2

an interesting article. I can't comment on why costs should increase in the land of 'infinite, free' bandwidth. The last time I saw a price quoted to 'buy' Internet bandwidth in the US (some years ago, so I imagine it has only got cheaper since then), for 10Mbps the quote was \$10 per Mbps. At that time the common wholesale price in Australia was \$400 per Mbps.

What I can say though, is that, whatever small decreases there are on the supply side, they are more than offset by increases in volume.

You can see on our total bandwidth graph for the last year:

versus the number of users connected to our LNS's in Sydney:

A 21% increase in connections, but a 45% increase in traffic.

In router processor terms, it is the volume of traffic that is more important by a factor of 3 - which for example would mean that an LNS router that can handle 10,000 connections where the average use is 20kbs, can only handle about 3,300 connections if the average use is 60kbs. So an increase in the traffic volume not only has the obvious effect of having to buy more bandwidth per user, but also means more/bigger routers at each part of the network - border, core and edge.

My guess would be that while US bandwidth is either 'free' (if you are a T1 carrier) or very cheap, the real cost driver is coming from the network engineering department with their incessant CAPEX requests for ever more and bigger routers.

Posted by Steve Waddington at 13:45

Wednesday, April 8, 2009

Entropy

There is an argument that the apparent paradox of life reversing entropy just, in fact, turns out to be a faster way of the universe achieving heat death.

I can tell you this though, it is amazing how fast entropy creeps in to what I once thought were robust, self sustaining systems and processes. I am not talking about anything serious, at least when caught early. Just little things like a broken auto-report here, or a process still being followed that is no longer relevant, or even a minor regular check that someone has forgotten to do for a while.

No one, or perhaps even no dozen of such minor 'breakdowns' would cause any problems in the normal running of things. Well, no obvious problems anyway. I think of our systems as having the same sort of fail safe redundancy as an aircraft - inherently robust.

But let these little bits of entropy add up, and one day the scenario is going to be worthy of an 'Air Crash Investigations' series.

About a month ago a series of three minor issues that came to my attention. It was a lucky coincidence, because if they had been spaced apart, I would have probably dealt with each one individually and not been alerted that the root cause was more systemic. The result was, not exactly alarm bells, but a feeling of unease relating broadly to helpdesk, fault tracking and internal escalations.

So I went a-hunting for what else could be wrong. And for the last month have set a task of finding at least one problem a day, and to have it resolved by the end of the day.

The first week it was no problem to find two or three things each day. Again, nothing serious, just anything as small as some text no longer linked on the web site that should be, or correcting an advisory email of some small typo.

Towards the end of week three, one a day was average, and this week I have only found one thing in the last two days. With some optimism I think I am nearing the end, as least as far as helpdesk processes go.

Combined with the web site suggestion box and a lot of work done by Larry to improve the resolution process, it seems to me a great deal of improvements have been made - which now that I say it seems somehow odd for a system that was working fine and doing everything it needed to do to start with. But there you go, there is always room for improvement and one must never rest on their laurels. and probably some other cliché's too.

Some of the trivialities have driven me to distraction, and I am sure the staff too, who I hassled to get things not just fixed up, but changed to not degrade again. After the fact it has been satisfying to get those little annoyances removed and see entropy reversed. One in the eye for the 2nd Law of Thermodynamics!

Posted by Steve Waddington at 02:29

Friday, April 3, 2009

No News is Good News

... or no news is better than evil news. And who wants evil news after all?

From a recent suggestion by a customer, I have been re-investigating the cost of an Undernet news feed, and the associated bandwidth and hardware required.

Five or six years ago I recall the bandwidth required for a 'full' news feed was in the order to 100Mbps, and for some reason, I expected that it might have increased a little, but not too much over that time. Shows how out of touch I am on that front. A full news feed, I am informed by two 'supplier' companies in the US, will require 500-600Mbps of dedicated bandwidth, bursting to 800Mbps. So effectively a 1Gbps feed.

I did a quick estimate of the hardware required too; about \$25,000 worth of servers. From previous experience, if you are going to run a news server, the people who want to use the service complain (bitterly) if there is less than 30 days retention. Retention of course meaning for the binary groups that hog the bandwidth, not the text groups. Which in turn means something like 100Tbytes of disk storage, or around \$50k.

A bit more than I expected to be sure. But providing a news groups service has always been the most expensive of all the 'free' services ISP's provided in the old days. Which is one of the three reasons have never done it in the past, the other two being:

1. There are plenty of really good, high retention and high availability fee for service news services available
2. Not that many people want it any way. At most I would say 5 in 1,000.

So, is it necessary to spend the \$75k on hardware, the \$150,000+ per month on bandwidth, plus the other associated infrastructure and OPEX costs?

Yes. Absolutely, if it is a service customers will pay for. For 'free' though, it is a tough sell. At best it might make 50 people using Exetel really happy, and attract perhaps one or two extra customers per month. In the mean time, the real costs associated with it do have to be paid to the various suppliers, and the cost has to be covered somehow. Which seems unfair for the other 99,950 customers who really don't care about a news feed to have to front up (albeit only a small cost each).

Anyway, I'll put in my \$25,000 share of the CAPEX. I'll just go and see if Annette wants to put in the other \$50,000.

If I were you, I wouldn't be holding my breath.

Posted by Steve Waddington at 12:57

Monday, March 9, 2009

FTTH

We looked at a wholesale product for Fibre to the Home this morning. The technology was interesting, but nothing new - which is good to see in many ways, because it points towards the service being solid and reliable, based on tried and true equipment. It was refreshing to find such a no-nonsense wholesale supplier.

I guess the downside is that coverage is limited to only those new estates and developments where their fibre runs. But what was also interesting was to see just what was possible with, essentially, 'off the shelf' technology. Something that puts the whole Telstra FTTN rigmarole in perspective (the implication being that the ONLY thing preventing cost effective 100Mbps speeds to the households of Australia is the monopoly of pit and conduit access).

The most interesting part however, was the cost it can be made available to the end user for.

We are currently looking at the business agreement, and still have to crunch the full cost for connection and supply. The early indication is that it should be possible to sell a 25Mbps fibre service for something less than \$40 per month, including some reasonable amount of data use.

Posted by Steve Waddington at 13:34

Monday, March 2, 2009

It Ain't Necessarily So

a ping time of near half a second
a ping time of near half a second
will stop VoIP from work'n, it's no use complain'n
if your ping time is near half a second

(sincere apologies to George and Ira)

So that is what 'everyone' will tell you. Most recently our VoIP PABX supplier when we were trying to fix a call drop out and quality problem at our Colombo office.

"The manufacturer says that VoIP wont work if the latency is higher than 120ms" we were told. And sure enough, the recommendation in the manual gave 120ms as the 'maximum permissible latency' recommended for VoIP.

Which we all know is total rubbish - meaning VoIP would have to be limited to only working within a single country or a geographic area not more than 5,000km between the furthest points. I can only imagine that the manufacturer quotes that figure (knowing full well it is well below what is possible), to prevent 'difficult' support problems (like ours perhaps).

The fact is, our PABX had been working just fine for over a year with no problem at all with return packet times in the 250-300ms range. What we really wanted to know was why, all of a sudden, the quality was degrading for some calls (but not others).

Obviously if it was only a problem for some calls, a general issue of latency was not going to be the cause. Never the less, our supplier stuck to their guns and could offer no further advice until we resolved the latency issue. Unfortunately, moving Sri Lanka and Australia closer together wasn't an option, and FTL quantum tunneling looks like it will be on our test bench for a while longer.

It just seems crazy that VoIP equipment manufacturers are still sticking to those type of latency figures, when for the last six years at least, the whole world has been using VoIP for low cost calling card and override calls around the globe. A call from Sydney to London is always going to be around 350ms, and when was the last time anyone noticed a quality problem? Unless you are paying the monopoly non-discount rate of some dollars per minute, the call is going to be VoIP.

To cut a long story short, we resolve the problem with a combination of bandwidth upgrade, a hardware accelerator card in the PABX to handle VoIP compression, and a software upgrade.

That was done a while ago, but calling home to Perth this morning from our Colombo office with this ping time:

```
C:\Windows\system32>ping 58.96.x.x
```

```
Pinging 58.96.x.x with 32 bytes of data:
```

```
Reply from 58.96.x.x: bytes=32 time=322ms TTL=238
```

```
Reply from 58.96.x.x: bytes=32 time=371ms TTL=238
```

```
Reply from 58.96.x.x: bytes=32 time=360ms TTL=238
```

```
Reply from 58.96.x.x: bytes=32 time=305ms TTL=238
```

```
Ping statistics for 58.96.x.x:
```

```
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
    Approximate round trip times in milli-seconds:
```

```
        Minimum = 305ms, Maximum = 371ms, Average = 339ms
```

The call quality MoS was a prefect '4'.

The end result of the less than helpful support from the VoIP PABX supplier, is that we are now in the final stages of development of our own Asterisk based PABX. The key to good quality on Asterisk (we have found) is to not skimp on

the hardware and especially not on the cost of buying commercial CODEC licenses.

Raymond and his team trialled the system with live support calls to Colombo last week, and the results were promising, with the same performance as the existing (fixed-up) PABX.

It looks like the all-up cost, sans development time, for the Asterisk box will come in around \$15,000. Or about 1/3 of the existing PABX cost. The biggest benefit for us though, will be the much closer integration we will be able to achieve with our backend systems, without having to pay the \$50,000+ for SDK's and API's.

Posted by Steve Waddington at 17:26

Sunday, March 1, 2009

This place is the A--hole of the world...

... may well be something Mr Truijlo would never say about Australia. Never the less, it would be a shame not mention that he was 'just passing through', as that old septic tank joke goes.

I guess he will be plopping off back to the States now. And from what I read in this AustralianIT story, his passing comes as a relief to the board.

Posted by Steve Waddington at 11:47

Tuesday, February 17, 2009

Measurement Lab Results

These are my results from M-Lab's sponsored 'Glasnost' test:

Image reduced to fit on page. Right click 'view image' to see full size.

An interesting result, since I know the NetEnforcer rules for my IP address would have been giving me a share of the p2p bandwidth only.

The indication is that a combination of properly applied 'time of day' rules, and PeerApp caching create an environment of no noticeable (to the end user) difference to P2P, yet still deliver significant bandwidth savings. Which our subjective testing showed was the case, but this provides an objective, external, independent qualification.

I'll run a few more of these over the next couple of days and see if the result holds.

Posted by Steve Waddington at 18:31

Thursday, February 12, 2009

Old Fashion

... with a twist, proving, even for my age I am old fashion, I am sure. (not that I am even that old).

In the very olden days it was considered 'proper' for a more junior employee to invite their more senior manager to dinner at their home. This gave the employee the opportunity to show their manager that their home life made them a 'solid' employee, that could make solid returns to the company on for any possible future promotion or pay increase. It also gave the manager the opportunity to gauge that the employee was a 'company man', on whom any promotion and pay increase would not be lost. because, as was well understood in those days, no man is an island, and in fact, no company considered that they were hiring an individual without taking their entire home situation into account.

You only need to look at the TV shows of the day to see how that was true, from shows like Occasional Wife, Bewitched and even The Flintstones, and just what comical situations of 'the boss over for dinner' would generate. Ever see that on a contemporary TV show?

In the seventies that changed. The rate of divorce had increased dramatically, and far fewer people, which is my guess only, thought that their wives/husbands were suitable people to introduce their managers to. Yet the imperative to managers was still the same - to gauge the stability and in some way the 'worth' of their staff to their company. Given that the very managers themselves were most likely in the middle of a divorce or some domestic dispute, the 'norm' became for the manager to take out the employee for lunch/dinner - with their spouses if they so wished.

That all changed again over the next two decades. People live much more insular lives now, and 'work' is far more disconnected from 'lifestyle' than it has ever been in history. So that sort of semi-work/semi-social interaction is universally far less used today. Most people can't be bothered. Almost no one considers the people they know at work, or their manager, to be of any relevance to their 'social' network (unless they were previously friends), that, presumably, they see contributing to their net worth as a person.

Yet what has really changed? I am not talking about the people that, in yesteryear where the 'stickflickers' of the world, never wanting a real career and happy just to fill in time between the weekends when they can do the things they really want. I mean the people who seek real advancement and improvement of their station in life. Have we really become such an egalitarian society that the bare minimum effort is now enough to ensure the career aspirations of all are met just by default?

Seems unlikely.

Of course, who can entertain at home now days? I am a firm believer in hospitality; Barbara, a great cook and hostess, and I, do our best to ensure any guest to our home has very nice time (and I have a very basic, but proven, strategy to overcome all the deficiencies of my very average education and poor manners, which is; ensure the beer/wine/scotch is good and flows freely). But my wife and my personal, traditional and archaic views aside, that is far from the expectation anyone has of any employee or manager today. Never the less, the impetus remains. So what does it actually mean when anything beyond the '7.5 work to rule' hours per day interaction is treated with a sense of total imposition? I can only guess:

- I hate the company I work for, I am only here until I find another job
- My partner and my mother in law really are witches. My mother in law often drops in unannounced, which is why sometimes I am such an ass - literally.
- I have no ambition, leave me alone, please
- I despise my staff/manager, I can't bear to spend any time with them apart from what my employer mandates I do
- I live in Australia, you owe me this job, just pay up and leave me the f___ alone

If it is indeed the case that the ability of any manager in a modern Australian business has become so pathetic that they can in fact not provide any benefit to their staff other than 'cracking the whip' with the illusion that they are somehow providing some benefit, then any and all of the above points are valid. If that is not the case, then for any individual seeing such a problem, the remedy should be self evident.

The hour is late, time for a Martini, a much more modern drink, that Larry, Darren, Sam and I can all appreciate.

Posted by Steve Waddington at 17:49

Tuesday, February 10, 2009

Change Carrier, Double Bandwidth Use

John and I were puzzling over this yesterday - In January we started to migrate out of contract services on the TW ADSL1 service to the AAPT ADSL2 service. Part of our plan being to reduce the total backhaul bandwidth with TW (our most expensive domestic bandwidth cost) to a level where we can retire one of those, very expensive, circuits.

While that has been underway, we have been looking at the expected bandwidth increase on the AAPT cross-connect circuit, which you can see here:

Our expectation was that we would see a proportional decrease in bandwidth across the two TWE circuits. But as you can see, that is not the case:

What has happened is that the a decrease of 50Mbps average use on the TW ADSL1 service has become a 160Mbps average use increase on the AAPT ADSL2 service.

You can see why 'why' is the question of the day.

My own ADSL service at home was one the the ones migrated in January, and so I have been able to run direct comparison tests between the two services. I am using a Cisco 827 router, which only supports ADSL1 encoding, so even on the new ADSL2 service, my router still syncs at an ADSL1 speed, which has been the same on both services, that is, a sync rate of 8Mbps.

Here are my results for service performance:

I must say I never minded 'only' getting 4Mbps download on ADSL1 - it was plenty fast enough for the mundane use I,(or any CTO I imagine) put it to. Nor do I or my family particularly notice the unexpected addition of 2Mbps capacity. But that is not the point. The point is, why _is_ there that extra capacity?

The only thing that has changed is the carrier. When I say 'only', of course that means everything from the DSLAM equipment, backhaul and carrier equipment all the way through to the Exetel POP. What I mean is, why should that make a difference?

As far as I could tell there was never any contention on the TW circuit. Download speeds would vary from time to time, but never by enough to indicate to me that there was any capacity problem. The latency is within 2 ms with either carrier, and the packet loss is the same (none).

Yet now we can see clearly that is is not just my own service, but there is a general doubling of the bandwidth being used across ALL of the 4,000 services moved away from TW. While my particular use of broadband means I don't notice/use the extra capacity the carrier change has bestowed, it is pretty clear that plenty of people do.

It is such a glaring difference.

We couldn't come up with any obvious reason for it at the time, but here are a few 'possibles':

- um
- er
- hmm

Ok, there really are not many option as to why it could be. All things being equal (which uniquely I can be sure of in these very specific circumstances), the only possible cause I can see is bandwidth restriction on the carrier network. I

am not saying bandwidth contention, I am saying restriction such that services like VoIP, ping and others that might lead to noticeable packet loss or degradation are not affected, but data downloads are constrained, by 50%.

Can that really be the case? It seems very strange to me if it is. While the results speak for themselves, it is hard to imagine any carrier with less need to restrict bandwidth than Telstra. Can anyone see any other possibility?

Posted by Steve Waddington at 09:33

Monday, February 9, 2009

"Optus Sucks Not True Claims IDC"

... was a headline of an article on the Smarthouse site that caught my eye, referring to coverage of the Optus 3G network.

The comment was apparently made by a journalist at Google press conference. Optus responded with 'no we don't, and here's the IDC report to prove it!'. Lucky those figures were on hand I guess.

Fortunately, as an all too frequent traveler, I am in a position to judge such a comment first hand, and need rely on no refutation by the carrier, or biased journalist's assertion. Further, there may be people who place more importance than me on hour by hour and sometimes minute by minute business use of the Internet. There may be people who get more than my average of 800-1,000, 50Mbytes, emails per day. And there may be companies that place more importance on Internet/Intranet access for all their staff, 24x7x365, than Exetel. But I can't think who, or in what circumstance that would be possible.

So I think I can safely claim, that as a 'power user' of the Internet, I have an extremely critical view of any roaming service that didn't deliver.

In fact, I have made previous blog comments on why WiFi doesn't work out, and why other proprietary wireless technologies also let down the traveling user. Something I was very happy to discover was not the case with HSPA.

The point I am trying to establish, is that when it comes to commenting on the 'suckiness' of a 3G network (in particular HSPA), I have, I believe, as much credibility as an IDC report, and much more than any journalist, or for that matter, any Optus exec. Based fully on directly on the first hand experience of the extreme use I put it to every day.

And what I can say is this; During my road trip to test out HSPA last year, and in the many places I have been since, testing and seeing just where/what the limits are - everywhere there is a Telstra signal there has also been an Optus signal (we have handsets on the two networks to test with). Everywhere there was an Optus signal, even as low as 'one bar', the data flow was sufficient for email and business web use. And wherever broadband was available, VoIP was usable with a MoS of 3.5 or higher.

Yes, there are certainly black spots in the coverage, in remote country areas. Like Gnowangerup for instance, or between Pemberton and Nornalup. But, my point is, there was also no Telstra coverage in those areas either.

The truth is that I am yet to find a place, coast to coast, where any reasonable person would expect a service be available that isn't, at speeds and availability that are, from my point of view as a real, heavy, business user, as good as it gets.

Posted by Steve Waddington at 10:28

Thursday, January 29, 2009

Welcome to Measurement Lab

Google have just released a set of on-line tools for measuring broadband performance:

<http://measurementlab.net/measurement-lab-tools.html>

Part of the hype around it is to 'out' ISP's that users suspect are throttling p2p or other protocols. It will be interesting to see, of those ISP's that deny (but we all know are) using bandwidth controls, just what their reactions will be.

I tried some of the tests this morning, but unfortunately most of them reported that the servers were overloaded and to 'try again later'. It will be interesting to see the results when the servers are available though. Over a fairly short time I think the tools will be very useful to establish baseline performance.

Posted by Steve Waddington at 11:17

Friday, January 16. 2009

BGP Security Concerns

From the article:

<http://www.networkworld.com/news/2009/011509-bgp.html?page=5>

I don't know what form such security is going to take, since the quoted development time is a couple of years. Nor is it apparent how it is going to be any better than the, already in place, BGP neighbor password directive.

I remember when a 'mistake' in configuration brought most traffic at PAIX to its knees. That was in 1999. It lasted for about 90 minutes, and I have never heard of another incident like it since. It was caused by slack filters on a tier 1 Telco router (I can't remember, but it may have been AT&T), combined with a very junior engineer (who should never have been left to do the work unsupervised) from a PAIX peer who made a basic config error.

The cause itself was removed very quickly (by a PAIX engineer ripping out the cables, shortly followed by some parts of the jnr engineers anatomy by all accounts). In those days, router processing power wasn't quite what it is now, and the rest of the down time was caused by routers running out of CPU as every router on the peering fabric tried to send and receive full routing tables from every other router.

We all learned a valuable lesson on that day - never rely on your BGP peer's filters. Or rather, those that didn't already know found out all about it as the lesson was rammed home.

So maybe, just maybe, such an event is a once in a decade occurrence. It obviously isn't more frequent than that, and I tend to think it may remain simply a 'one off' aberration.

The claim that extraordinary BGP security is needed to prevent accident mis-configuration simply doesn't stand up.

By the same token, and for the same reason, it is also hard to see how deliberate hijacking of BGP could cause any major disruption. It would require widespread insertion of bogus data at many, many points to have even country wide effect. Even in a country as small as Australia, at least 20 separate data centres/cable plants/backbone trunks would need to be compromised, at the same time to cause major national disruption. And just how do you hack into a fibre trunk without being noticed anyway?

Anything less in the 'attack' is going to be subject to the same 'rip out the bad wires' fix of the PAIX incident. It is going to affect a few people, for a short time, but the Internet is inherently robust to such disruptions.

I would think a far, far greater threat, and way, way easier to do, would be to nab a Network Admin, who knows the enable passwords, on his way back from the pub on a Friday night.

From my own point of view, I would like to think I would hold out and 'take one for my country'. But the reality is far more likely to be that the first touch of the blade to my pinkie finger, and the router password are theirs.

Hard to see how any amount of BGP encryption is going to overcome that.

Posted by Steve Waddington at 15:38

Saturday, January 10, 2009

The truth is out

... and there can be no denying it. I mentioned a while ago that the phrase 'worst [thing to happen in the economy] in 25 years' was used, and wondered how long before that 'worst thing' currently would start to be compared with the great depression. This morning on the radio I heard 'Worst job losses recorded in the US since World War II'. So there you go, we are on the very doorstep of the great depression. And the news bulletin went on to add the even less cheery news of 'forecast of worse still to come'.

I still have memories of the good old days though. When was it now? September, October? When our fearless leader in Canberra was telling us that everything would be fine, no need to worry. I wonder what he meant? Maybe what he was really saying was 'Everything will be fine (for me, I have a secure well paying job, at least for the next three years)'.

Anyway, it does make me wonder - you would think that if anyone was well briefed on the economic outlook it would be the government of the day. They have the best access to the most information of anyone. So what can be implied from the 'everything will be fine' type statements?

I can only see two possibilities:

1. A condescending remark from Fearless Leader to his 'children' to put our minds at ease and spare us, at least in the short term, from worrying about what is to come. A lie if you like, albeit a white one, albeit patronizing.
2. He had no idea either.

Bit of a worry either way really.

Posted by Steve Waddington at 11:12

Friday, January 2, 2009

Remote Site HSPA II

We have now set up the Yagi antenna. Although not yet optimally positioned, the signal strength has increased from 19% to 29% (as reported by the Netcomm modem).

Lining up the Yagi with the Mt Barker base station.

'Fine Tuning' with a 10" shifter.

The improved reception increased the download speed as well - six tests down via Speedtest.net returned results from 1.7Mbps to 2.3Mbps download speeds and upload speeds between 250kbps and 310kbps.

While there was noticeable improvement with the higher gain antenna, it was still less than I was expecting. But a little bit of extra hight should improve it significantly.

Despite the relatively low signal strength readings, the circuit itself has been rock solid. It has in fact been performing better than some of the 'in town' locations I have been testing from.

Posted by Steve Waddington at 11:33

Wednesday, December 31, 2008

Remote Site HSPA

I installed the Exetel 'Mark 1' HSPA kit at my friends Bill and Sharon's farm block this morning. Currently it is using the standard mobile car whip antenna that was already in place, as a little work needs to be done before we can mount the Yagi antenna.

The signal strength from the whip antenna showed as 19% on the Netcomm GSM modem stats, however that equated to 1.3Mbps download and 100kbps upload speed, zero packet loss and 250ms ping latency to Sydney. These are the results from two NetSpeed tests I ran:

<http://www.speedtest.net/result/381569601.png>
<http://www.speedtest.net/result/381571339.png>

This is the current antenna setup. We will replace the whip antennas with an 18db gain Yagi.

The previous setup required two antennas (and the two GSM accounts). The Exetel setup needs only one.

This is the installed kit:

1 x Netcomm N3G002W, 1 x Netcomm V210P and a \$22 handset from Dick Smith, provides broadband Internet and VoIP.

This is the shed from the front gate of the farm block. Anyone good at 'Where's Wally' will have no problem spotting the shed behind the trees in the middle of the picture.

In the middle of a 40 acre block.

Even with 19% signal strength, the broadband performance was as good as an average ADSL1 service. We made several calls to local Albany numbers using the VoIP service, with MoS ranging from 3.6 to 3.9.

Posted by Steve Waddington at 13:47

Monday, December 29, 2008

HSPA - Brilliant!

We have just arrived in Albany for a bit of a break for the family, and so I can test the country-area HSPA kit we have put together. I will have some pictures of it tomorrow, fully deployed, but for now it is still packed in the car.

But what I can say right now is just how very, very convenient HSPA is. When I did the road trip last month, Albany and the lower great southern only had dial up speeds via HSPA. Since then, a couple of people have reported that it has been upgraded to full 2.x Mbps in late Novembers, so I had an idea what to expect.

Having an idea of what to expect is quite different from seeing it first hand though. Here I am sitting in a holiday unit with FULL 2.254Mbps broadband access, near perfect VoIP, and 100% access to our management reports and internal systems through my fixed, secured IP address.

Of course, I don't need to pay \$5 a month extra for Internet access via HSPA - I could just get \$10 pre-pay dial up access from the news agent, then it would only cost me the \$10 plus the 50c flag fall for a local call from the hotel room phone.

Or, I could use the resort WiFi for only \$22 per 24 hours.

While I have used, and appreciated both of those methods in the past, I am glad to say they are now in the past, and I shudder at the thought of ever having to go back to such inconvenience and cost.

For any financial controller or business operator with staff on the road, it is also hard to see how they could pass up the trivial extra mobile bill cost for the certain cost savings on remote Internet access.

In actual fact, just this four day trip away alone will pay for the HSPA cost for a year. How did I ever get by without it? It's so brilliant you could paint it black and white and call it a pint of Guinness.

Posted by Steve Waddington at 18:55

Friday, December 26. 2008

And so that was Christmas....

... and what have you done? Misquoted from that great John Lennon song. Apart from the really, really big network changes we did, by necessity, in a few early hours; if you asked me on any day what was different, or even any week what was different to last week, the answer would probably be 'same, same'. Yet looking back, it has been another amazing year in just about every respect. The list is long, though for me the highlights have been:

- Business grade services sales doubled over the year
- In the face of the free give away war between the two largest ISP's, Exetel sales remained on track
- Everyone employed by Exetel received a pay rise over the year
- Everyone received a 'Christmas bonus'
- POP's established in all mainland states
- Capacity of the Sydney POP increased five-fold
- Capacity of the Exetel network increased ten-fold
- developed a robust VoIP product
- Completion of the GURUS 'alpha' phase
- Integrated two new broadband supplier/products

But, the thing I can say I find most satisfying is that we now have 12 more highly competent, highly skilled engineers in network engineering, system administration and software development. Not people we hired externally, but people who, starting in helpdesk support, have been able to move up within the company and begin to fulfill the potential for which they were first hired.

Doubly satisfying is that, as a small business, we have been able to grow and accommodate the career advancement of our people over the years.

Anyway, that's enough looking back - another year over. Over the next few days I will complete mapping out the likely network and systems changes we will need in 09, starting with the country HSPA product. I have the 'mark 1' build kit sitting next to me, and will take the opportunity while heading south for New Years to test it out to its limits.

- and a new one just begun.

Posted by Steve Waddington at 12:45

Thursday, December 18, 2008

UB2000 back in full operation

The replacement switch for the UB2000 system arrived late in the afternoon on Tuesday. And during the day we had had time to consider how we would re-integrate it.

The emergency work-around we put in place Tuesday morning was working fine, and there was no stress on the network, no observable problem, and most importantly, no customer reports of anything being adversely impacted.

There was no technical barrier to just reversing the change and putting the new switch in as a drop in replacement for the broken one.

However, and there is always an however..

Doing that would mean a necessary break in service for a few minutes while we physically disconnected and reconnected the Ethernet cables from one device to the other. Maybe 5-10 minutes service disruption at most.

If there had been problems evident, or some pressing issue as a result of the workaround, then the decision would be easy and we would have made the change same day. (there is a possible window for such changes between 5pm and 6pm - one of the lowest traffic ebbs of the day. Hard to use the Internet when you are on the way home from work I guess).

But since everything seemed fine as it was, the more sensible time for the change back was the usual 7am to 8am maintenance window period on the next day.

We still didn't know exactly what effect not having the cache would have during the evening peak period, though every indication was that we could a) control any contention caused by P2P though more aggressive application of NetEnforcer rules, and b) wasn't going to be severe in any event.

Which is exactly how it turned out.

A couple of minor problems were encountered on Wednesday morning during the swap back. Mainly to do with an initial ping test failing, requiring an arp flush, but it took a little while to figure out. And with a missed line in the route-map code that directs high port traffic to the cache.

Since the plan was to swap one circuit first, test, and then swap the second, the result was we ran out of maintenance window time satisfying ourselves that all was good with the first circuit.

As an aside; a rule of thumb for maintenance windows is to allow 1 hour per expected 10 minutes of service interruption, plus 15 minutes either side safety margin. Since we had expected 5 minutes disruption per circuit, in retrospect, we should have given ourselves a 1 1/2 hour window rather than 1 hour.

Never the less, the change plan catered for such an event, as each circuit could be migrated independently.

Rather than push the window, the decision was made to complete the 2nd circuit the next day. Again with during the 7am to 8am period.

Everything went fine. There was some initial 1% packet loss monitored - a combination of cache-burst and high router load due to BGP sessions re-establishing. A few adjustments were made to the load balancing and it was done.

One bit of good news - PeerApp have kindly provided a 3rd switch which we will keep right next to the other two as a hot swap spare.

Posted by Steve Waddington at 11:44

Tuesday, December 16, 2008

He's Dead Jim...

What is it about power supplies? A few months ago a power supply on a 7301 went out, and early this morning the same thing happened to one of our suppliers 3650 switches. Dead, dee ee dee dead.

The problem with that particular switch, being the front end of the PeerApp UB2000 cache, is both GE circuits feeding NSW and other states with Internet traffic are connected to it. So the switch outage also took those two circuits, and hence all Internet access into our Sydney POPs, except for peering, off line.

The switch went out at at 4:30am this morning. About half an hour later the problem had been positively isolated to the particular switch, and since there was no remote access it (it being dead and all), Jing, one of the duty engineers was dispatched to the co-lo site.

Since I was in Sydney, and up anyway (funny how that happens), I decided to go to the POP too in case I could be of any assistance. I arrived just after Jing, who has already confirmed that the switch was completely out.

Our options from that point were 3:

1. Get a replacement switch via the supplier
2. Use the one remaining switch and reconfigure our own 6500 to take the place of the failed switch
3. Bypass the UB2000 altogether and connect the Internet supplier circuits directly to the 6500

What you don't want to do in a situation like this is panic and start changing things without fully considering the ramifications. Even though the time pressure to get the service fixed is intense, it is far better to take a few minutes and think/talk through the options and proposed changes than to just barge ahead and make whatever seems to be the best change at the time.

William was on the phone to the supplier, but it didn't look like option 1 was going to pan out in any reasonable time frame (ie immediately).

Option 2 was at face value the best choice, because it offered a return to full functionality, including the UB2000 cache. However; it meant some fairly serious route-map code being added to the 6500, along with a drastic topology change. Adding the code itself would be trivial, since it was already written for the 3650, but what we couldn't be sure of is how the 6500 would perform when pressed into that role. Probably it would be ok, but if it wasn't, then we would lose the whole network until we rolled back.

Option 3 meant losing the UB2000 cache, but the change could be accommodated with just a few config changes and two cable moves. Most importantly, there was zero chance of anything unexpected happening. Option 3 it was then.

Taking it steady, to be sure, to be sure, we tentatively re-routed one circuit. Kaplah! and the bgp session re-established itself before I had time to type in 'sh ip bgp sum'.

There were no obvious problems with that circuit change, so after 5 minutes we made the second change. Success also, and we were pleased to watch traffic levels return to time-of-day normal over the next few minutes.

The NSW network is now running without the UB2000 cache, so it will be of some interest to see just what impact that might have on traffic during the day. I expect it will not be great - maybe a slight degradation in some P2P performance. We will have to wait and see.

We should know the delivery time of the replacement switch within the next hour. Depending on how the traffic looks, we will probably schedule the replacement early tomorrow morning, since it will require about a 5 minute outage to change the 6500 config back and re-re-route the cables.

Posted by Steve Waddington at 09:16

Tuesday, December 9, 2008

The value of money

My Grandmother knew the value of money. The only thing she didn't pay cash for was her house, and she scrimped and saved and sewed and knitted, and somehow, on a combined income of 20 pounds a week, managed to pay off the 1,500 pound mortgage in five years.

Like my parents, she never owned a credit card, never took out any form of loan for anything (other than the one house she bought), and had a deep suspicion of any form of usury. 'Neither a borrower nor a lender be' was a quote I heard more than once, growing up.

In my grandmothers day, any sort of debt was to be avoided, and 'loans' (the term 'credit' was either unheard of, or about equal to another 'c' word, and never used) were for those without the moral fortitude to save up for the things they wanted.

She was a woman who lived through hard times - born before WWI, married at the start of the depression, her first married years were in a dirt floor slab hut in Jarrahdale (where she gave birth to my Mother). She was a tough, pragmatic woman with her views, and wisdom, shaped by the tough times she lived through.

I never once heard her complain about what she had been through, nor ever criticize the, which to her must have looked like, very easy and soft life her grandchildren had. There were some things that she just would not talk about, and she always had a lot of time for the Salvation Army, so my guess is that she went through some events that she did not care to burden anyone else with.

Without doubt, she knew the value of money, because for most of her life it was so scarce. She didn't hoard it, but she was very careful with it, and never spent more of it than she had, with a mindful reserve for the unexpected.

The values and beliefs of that generation seemed to rub off on my parents, who also have never had a credit card, although I know have had loans for cars, as well as houses (paid off mostly in the normal loan time). But sadly, the benefit of that wisdom didn't quite seem to make it to me, and like most people in this age, I have the requisite number of loans, credit cards etc.

So here we are, most of the world is officially in recession. And the 'worst fall/drop/time since...' keeps being pushed back further and further - 34 years I think was what I heard yesterday as the worst drop in some share price index.

I have no idea what the governments splurge of ten billion dollars will do. It is hard to see how, even if it all were spent on tv's, washing machines or Christmas presents, how it will have any effect past the 1 month it all gets spent in. It will probably deplete some of the warehoused or stockpiled goods that have built up from money not spent in previous months. But will it keep even one factory in China from closing, or one iron ore shipment from being refused?

But, the poor people have their free hand out. Urged to spend it on 'stuff' (because, heaven knows, none of us have too much stuff already). Again, I have no idea of the economic theory behind it (yes, I understand the concept of fiscal and monetary policy, and the three main theories of what to do to correct 'negative growth'), but to just treat money like that, throw it around, as if it is worthless. It seems fundamentally wrong to me. And fundamentally irresponsible for those who hold the purse strings of the nations wealth, to treat it like that. (If money runs out, just print more - now where have I heard that before?)

Maybe some of that elder generation wisdom did rub off on me after all.

I know what my grandmother would have done with the handout (which she would pragmatically accept without comment). A bottle of good brandy, bacon on the table (rather than pigs trotters or ox tongue), and the rest in Shell Oil shares.

Imagine if everyone did that - \$10bn pumped back into the Australian share market in blue chip shares. Surely that would be the gift that keeps on giving.

Posted by Steve Waddington at 09:24

Monday, December 8, 2008

Plumbing New Depths

Is it just me? Does anyone else see this as an all time low for the behavior of this carrier?

1. Give us \$10bn or your national communications network will be stuffed. And of course we will hold you to ransom on exclusivity.
2. We can't be bothered putting in a proposal for a mere \$10bn, but here is 10 pages so show the utter contempt we hold you in. What? \$1bn per page isn't enough!
3. Well F--- you, we don't need your stinking money after all, and never did, but thanks for letting us screw the whole process up so now one else can get it either.

For some reason I have this 1960's western image in my head:

Calvera: Hey amigo, we can get \$10bn from these gringo's.

Bandit #1: \$10bn, that is a lot of peso's amigo

Bandit #2: (whispering to bandit #1) Shhh, only 10 billion pesos now hombre.

Calvera: Si compadre but it gets better, they give us \$10bn, and we also claw back our old monopoly

Bandit #2 Ai, Ai, Ai, but what do we risk?

Calvera: That is the best part amigo, there is no risk. We don't need the money, but if they don't give it to us, we can make sure no one else gets it either, so we still keep our monopoly and get the money from the campesinos in the long run.

My question is, where is Yul Brynner and James Coburn when you need them? (Ok, for the purists, maybe the Akira Kurosawa original was better, but no one can dispute that Elmer Bernstein's score is unsurpassed)

Posted by Steve Waddington at 10:38

Thursday, December 4, 2008

Pocket Money II - the Agent: Lite Program

Shady has completed the last of the web page tasks and the link is now added to the User Facilities area. Just in time for the holidays.

Now, any Exetel customer can register themselves, or someone (like their teenage children in my case), as a 'Lite Agent' for Exetel. This is the introduction page for the program:

Exetel Lite Agent Program

Who the Program is For

The Exetel LA program is exclusively for customers of Exetel. Its concept is to allow Exetel customers, in particular their children, to be able to earn extra 'pocket money' through recommending the Exetel ADSL service to friends, family and neighbors.

The LA program is essentially an extension of the recommendation credit that has always, and still is, available to Exetel customers.

Commission and Facilities

An Exetel Lite Agent will receive a once off payment of \$30 for each new customer who enters the LA's code on the order form, that then connects to Exetel. The payment to be made to the LA is calculated at the end of each month and paid by direct deposit into the nominated account by the 5th of the next month.

The customer signing up will receive a once of \$15 credit (by entering the LA's code on the order form) on their first months bill.

Lite Agents are 'recommender's' of the Exetel service only. That is; they are not affiliated with Exetel in any way and do not represent Exetel in any capacity. Lite Agents do not resell the Exetel service and receive only a once off commission for recommending the service to someone who then buys it.

An Exetel Lite Agent has access to a password protected LA Facilities web page where they can:

- * See a record of who has signed up
- * Access records of payments made
- * Access and print Exetel brochures

Becoming an Exetel Lite Agent

Only the Exetel account owner, or authorized adult can sign up for the LA program. They can however nominate one or more of their family as the LA contact, making a separate application for each person.

The sign up form is here:

https://www.exetel.com.au/members/ma_form.php

and requires the Exetel user facilities username and password to access.

Once the application has been completed a confirmation email will be sent, followed by an acceptance email within the next two days, which will contain a username and password to access the LA facilities area.

At this point the LA should also register on the Agent Forum:

<http://forum.agent.exetel.com.au/>

The Agent Forum site is the main support area for Agents and Lite Agents. Brochures and Support Material

The LA facilities area has brochures and other support material that can be downloaded and printed. The LA should ensure their agent code is in the appropriate area of the brochures, which can then be distributed.

Please note: using an inkjet printer to print the brochures is not recommended, unless only a very few are being printed. Inkjet cartridges are expensive, sometimes costing 50c or more per page. Business stationary suppliers such as Office Works, or even the local news agent are a much more cost effective method to print 10+ brochures.

Any comments or suggestions on how to improve the program will be welcome - email me directly - stevew (at) exetel.com.au or Shady, the program coordinator - shadyd (at) exetel.com.au

Posted by Steve Waddington at 20:24

Wednesday, November 26, 2008

Communication Breakdown

It's always the same... (Page/Jones/Bonham)

Will the AFACT case really 'shut the Internet down' as APC claims it might?

No. And for an obvious reason. What an ISP is required to do under Australian law (an opinion Exetel has held since it started, and recently backed by an opinion of an SC) is not all that hard for the ISP to do, and we (Exetel) have in fact been doing that for the last four years.

Data still seems to flow through our network and our 53,000 odd ADSL customers still somehow get sufficient access to everything they want to consider the service reasonable enough to keep subscribing.

Amazingly it seems it is possible to do the right thing, and still exist as an ISP.

Why would not every ISP, particularly the larger ones, who are presumably run by sensible business people, not take those same actions that seemed obvious, and right, to us?

I posit the view that is it because, essentially, those businesses have relied on the advice of their technical/sysadmin/engineering people, and as a result, ego's rather than common sense has been the rule of the day.

I am sure many will be familiar with the sadly all too common attitude of "My knowledge of this server/router/program is vastly superior to yours, and therefore I am vastly superior to you, and will treat anything you tell me with the utter contempt it deserves". And woe betide anyone who attempts to tell such a person what to do - like a legal notice to take action, for example.

There is a great generosity of spirit imbued into the Internet by its early pioneers - Postel, Cerf, Thomlinson, Kilgour and many others. Carried through to today with fantastically valuable (to the world) projects such as Wikipedia. Yet I have not observed that same spirit in abundance, or at all, in the Australian community of Internet technocrats. Just look at the petty sniping that goes on in forums such as Whirlpool if you want proof of that.

Worldwide CTO's and CIO's are starting to 'come out of the computer room' and take their place as valuable members of boards of companies - a rightful place for them to be as well rounded business people with expertise in a particular area, balanced with the business skills in other areas other individuals will provide.

But for Australian ISP's, I would say the last voices that should be listened to on any serious, mostly non technical, matter are those very people on whom they currently seem to rely.

My guess is that a little humility would have gone a long way to saving a lot of money in this particular case.

Posted by Steve Waddington at 16:07

Monday, November 24, 2008

Good Times, Bad Times

You know I've had my share... (Page/Jones/Bonham)

I was chatting the other day with my eldest son and some of his friends, who, since they had just completed their high school graduation, were expressing interest in what the future would hold, given this current economic climate. Immediately, I suggested to them, it was not going to be much of a problem, since for at least the next three years they would all be at Uni. By the end of that time, who really knows, but it is quite likely the world economy will be back on the upswing.

But even if it isn't, then what?

Only two things really, as I see it. One; things will get better, and every stock share and every house price will, within the next 10 years, be double what it was at the peak of the last market peak (ie, the All Ords will be 13,000). Ten years after that, it will be double again.

Or, two; civilization as we know it will end.

The last time that happened - that is, a reset to zero where all industry, fiat money and national wealth is worthless, for a very long period of time - was the collapse of the Roman Empire (I can't really think of another time in western history. Even though states have failed since then, western civilization as a whole hasn't gone backwards since about 800AD).

So, either things will get better in the next few years, or we have the apocalypse. Now there's a cheery thought.

The question is whether to buy gold or shares and property (assuming anyone has any cash to do that with). If you think the end is nigh, then gold is the only choice. And perhaps some guns, and a lot of ammo. According to one documentary on the future, there could also be a market for V8 interceptors and Kawasaki motorbikes.

What does any of this have to do with network engineering? Well, the real indicator of that will be to look at Cisco's Q1 2009 results (Q4 2008 was up on both Q3 2008 and Q4 2007). I can't see public or private companies committing money to new builds and more infrastructure. What it will come down to is how much of the G20 'rescue package' spending makes its way into Internet infrastructure spending.

I will have to think about what that might really mean for Exetel. In the mean time, I have no plans to change my plans - which is to have no plan to spend anything not absolutely necessary for a while. Not that that's much of a change to any plan in the past anyway, despite what my partners think :)

Posted by Steve Waddington at 11:40

Saturday, November 22, 2008

Stupid, Stupid People II

From this article:

Dalby also said it was nonsensical sending notices to customers when there was no direct link between an IP address and a person. "When we get notification from AFACT all they do is provide an IP address. It's not the equivalent of a phone number," he said. "There's no pointer that this IP address is Steve Dalby."

How uninformed, or what dysfunction within the organization would there have to be for a, very senior, representative of a public company to knowingly make such a statement?

I mean, he is not a technical person, so his comment can only be based on either his request for, or what was given by, whatever briefing he had from the CTO or equivalent, right?

Is he saying they can't match an IP address to an end user, or that they don't know the address of their end users? If that were the case, then they couldn't bill for traffic, because there would be no way to match the data downloaded with the user account. Yeah, right. That will stand for 0.1 seconds of cross examination. And it is the case, shouldn't the shareholders be concerned?

To make a statement to the press is one thing. I wonder how he will go if he has to say the same thing under oath (and threat of perjury).

Good luck with that fella's.

Posted by Steve Waddington at 07:55

Friday, November 21, 2008

Don't call us...

... and we wont call you.

It is the perfect relationship between supplier and customer. Supplier provides a service that always works, and the customer pays their money and never calls.

Sure, for an ISP the service isn't going to work perfectly all the time. But it does for most of the time for most people, and in those few times, or for those few people for whom it doesn't work, a simple trouble shooting guide and an IVR can take care of 95% of the problems.

Why then, is it taken for granted (a startup premise of Exetel is that we didn't, by the way) that there has to be a call centre and helpdesk staff to answer customer calls?

Since even before Exetel was started, John's mandate, and our joint desire, was to reduce the support calls we would ever need to answer. From the initial premise of 'no call is to take more than 3 minutes', we have put in place, and I have measured the effects of, actions that would let us reduce the number of people we needed to 'answer the phone'.

I think we have been somewhat effective in that. I mean 14 support staff for 90,000 odd customers is probably right up there as a ratio right?

But I think we can do better.

When I analyze all the reasons someone might call for support, rather than using one of the many, many other (and almost always faster) methods of fixing their problem, all of them are something like:

- end user didn't read the manual
- end user set up modem incorrectly
- end user changed setting
- end user doesn't like using IVR
- end user was lonely and wanted someone to talk to

So you can see what is common to all - the end user is always the one that cops the 'blame' for the call. Someone could, possibly even successfully, make the argument that 'you can't change the way customers think, or their habits, so we will always need to be able to answer support calls'. I know that argument is not true, yet, even though our ratio of support people to customers is good, we have only ever increased the number of people we have to answer phones.

Let us suppose for a moment that the reason for the calls has nothing to do with the end user and remove them as the cause. Could we then reasonably expect that we would only get calls when there was some network fault (given that customer hardware is replaced by the manufacturer, and provisioning issues are not support issues, but in any event automatically handled)?

If that were the case, then, since in an average month there are less than 10 fault minutes, we should only expect to get support calls for at most, half an hour.

Obviously that isn't the case, and calls come in all through the day and month regardless.

Therefore, while the network faults may cause an increase in support calls, they are not to blame for the average call volumes.

If it is not the end user, and we have established it isn't some hardware or network fault, what are we left with?

Since there are only three components, end user, network and support, the revelation we are forced to come to is that It can only be the support staff that cause the calls.

Rubbish! You don't know what you are taking about!! Stone the Heretic!!!

But why not? Put aside the thought that a support person can't influence and end user to make a support call for a moment. If you are a support person whose job is to answer phones, is it not better for you if there are more calls? If you are a manager of a support centre, it is in your interest if there are more calls, or less? And what are the KPI's an outsourced call centre would use to claim payment and bonuses on? 'Number of calls answered' would have to figure prominently.

The fact is, no one on the supply side of answering calls has any interest in having less calls to answer, and it is almost always to their direct benefit to have more calls.

But what if that wasn't the case, and the support person had a direct incentive to take less calls and spend less time on the phone? It is certainly in a companies best interest to have less calls. And in fact we have in many ways 'incentivised' people to reduce support calls, as the method by which they can progress within the company and their career to better and more interesting things.

Also, a number of initiatives, such as the EEE certification program, white paper wiki and technical briefings have all been aimed at putting in place the mechanisms whereby people can 'bootstrap' their careers past helpdesk phone support, and their managers too for that matter.

Maybe 6,500:1 is a good ratio for support staff to customers. But I think I will set a goal, without detriment to any other business metric, of 20,000:1 for next year.

Posted by Steve Waddington at 16:32

Saturday, November 15. 2008

Shutting down botnets

I watched the movie 'Untraceable' a little while ago. I thought they did a reasonable job with the dramatized use of technology. However implausible it would be to do in real life, there were no glaring technical errors (at least that I picked up 'going with the flow' of the movie).

What I did think though, was that Agent March gave up way, way too easily in shutting down the murderer's botnet sites that were used to propagate his web site. It took her about two mouse clicks to black hole an individual site, but then the devious criminal mastermind's site would reappear in a few seconds hosted on some other compromised system. At that point, the Agent March just gave up, because apparently no matter how many times she black holed a particular IP address, there would always be another the murderer could use.

But if you think about it, just how many computers can there be that any particular hacker has access to? 100, 1,000, 10,000? Probably closer to 100 than 10,000 I would think. So, in the movie, with maybe no more than 500 mouse clicks, black holing each new site host as it came up, would have saved a life and stopped the murderer in his tracks. Yet that was too much trouble, and so another three or four people had to die (horribly).

I am not saying this to criticize the movie, I fully understand the necessary 'fudging' that plot devices need. But think about it irl, even if there were 10,000 compromised systems, it is quite within the power of any network operator to block them all, even manually if they had to.

We have just put such a block in place ourselves for one dozen /32 IP addresses that we have logged attempting to gain unauthorized access to our systems. We have done so in the past, usually where we see related attempts to gain access to some core device from a single IP address, but this is the first time the same attempt has been logged from a number of different IP addresses - clearly the same attempt, but from such geographically diverse locations it can only be a botnet.

While we have always done this manually, since the incidences have been relatively few, this time we are also putting in place a script that will check logs and then put in place an auto block for such attempts. I am not going to say what the script looks for in the logs, but I can say that I think there is very, very little chance of getting a false positive (that is, blocking something that it is not meant to).

I am pretty sure the script can act faster than a botnet would be able to, and since our theoretical limit of what it can block is exactly equal to the size of the Internet, it is very hard to see how a botnet operator could marshal enough compromised systems to defeat it. I will concede it would make a very boring movie though.

Posted by Steve Waddington at 10:37

Friday, November 14. 2008

Stupid, Stupid People

'Never underestimate the power of human stupidity', one of the all time great quotes by Robert A. Heinlein. But after reading this story, I think also that the quote itself should never be underestimated.

Despite everyone telling her that she was being scammed, she kept sending of the money anyway - for years.

Still, who are we to criticize. Collectively we have managed to elect a government hell bent on spending millions for years and years to come on a scheme that everyone who knows anything about how these things work is telling them it will be a complete waste (or my own lonely view, that even though it is pointless, if it is going to happen, it can never the less be done at zero cost).

But surely they (collectively lets call them The Ministry of Information) are not acting out of sheer ignorance. There must be someone, or some group of consultants advising them that this is the right way to go.

Just who are these 'advisors', and what possible credentials could they have to be able to give such clearly wrong advice?

I think the real question however has to be, who is going to gain from this?

Posted by Steve Waddington at 09:35

Thursday, November 13, 2008

WA Speeds

An issue was discovered with the IP bandwidth to WA yesterday. It was the 'good old' duplex problem that sometimes happens on a 10/100 Ethernet circuit when both the router and the switch interface are set to 'auto' to negotiate speed and duplex. Sometimes they will set themselves to half duplex, when both can support full. But sometimes also one will set itself to half while the other shows full.

At traffic levels of less than about 40Mbps it usually doesn't cause a problem, but higher than that excessive collisions, and consequently packet loss, start to occur. It is not so noticeable with small packets, but large, 1500 byte, packets can see packet loss of over 20%.

What actually happened in the Perth router's case was that initially both the router and the switch synced in at 100/full. Then sometime in the last 36 hours the router port at least 'flipped' to 100/half. We can't see any logged reason for that to happen, however it may have been a port reset on the switch in the early hours of the morning (being a supplier switch, we don't have access to it to check) that did it.

For whatever reason, we ended up with the port in half duplex mode, and as the traffic increased towards 50Mbps, large packets - typically those that would build images on web pages, began to be dropped. Anyway, once spotted, the problem was easy to fix. Both the router port and the supplier switch are now locked in 100/full mode.

Now, the really interesting thing about WA speeds.

A week ago I put in a router to terminate a direct IP feed from Perth, rather than taking IP traffic from Sydney. From that time, all WA users have been accessing the Internet directly via the Perth supplier, which means two things:

1. The WA traffic is not being controlled by the Sydney NetEnforcer
2. There is no Sydney cache access for P2P

With no P2P traffic control, yet with plenty of spare (well connected via a major T1 supplier) bandwidth, I would not have thought that there would be much noticeable difference in data transfer rates.

Wrong.

Downloading the latest Ubuntu ISO's with PeerApp cache access, I would most of the time get 4-4.5Mbps download speed on my 8Mbps ADSL1 circuit.

However, downloading the same ISO's without cache access (with great pings, no contention on the supplier bandwidth etc), the best I could manage was 1.5Mbps.

It is interesting to see just what a difference the PeerApp cache makes. The reason I put it down to is that the cache is closer, latency wise, than the majority of p2p peers. Even so, I would not have expected it to make a threefold performance difference.

Now what I need to figure out is how to get the best of both worlds so we use the supplier bandwidth and have cache access at the same time in WA.

Posted by Steve Waddington at 09:05

Wednesday, November 12, 2008

The Internet Newspeak

What a fallacy. That we have a government that thinks it can surpass the absolute total thought control that a genius like George Orwell could imagine.

The best those middle weight intellectuals (I use the term with extreme generosity) can hope to really achieve is the bureaucratic nightmare of Terry Gilliam's Brazil.

Wow I look forward to living in that Utopia - not.

But close our eyes and cover our ears. Stamp our feet and yell to drown out the voices that gainsay us, and we will have our own Hy-Brasil. The perfect utopia, and never shall a drop of blood be split.

Read this and make up your own mind:

http://www.dbcde.gov.au/communications_for_business/funding_programs__and__support/isp_filtering_live_pilot

Well, apparently, few people can actually do that - make up their own mind with an informed opinion that is.

How else would you explain the total extravagance of the expenditure proposed. \$140M. That is a fine use for my taxes - feeding those parasite consultants and ISP's that see 'government sponsored' as an excuse to dip their snouts in the tax payer funded trough.

So sure, it may well not cost the ISP anything to actually implement such a filter, because the government - actually the tax payer - will subsidize it. As a shareholder in an ISP I guess I can be happy about that. However, as a tax payer, I think the appropriate response to a pointless waste of money like this is outrage, because:

- a) as I have pointed out, such a 'blacklist' is incapable of stopping anything
- b) even if it were, anyone able to spend a 2 hours or less doing basic research would realize that such a proposal would not work - since it has not worked in 15 years since the first standard for RBL's was implemented.
- c) even if that did work, it is completely foiled by any sort of encryption commonly used on the internet today - for banking, on-line purchasing, or any thing else at all other than clear text.
- d) That stopping 'accidental' access to, alleged illicit web sites is exactly the same as trying to stop a shopper 'accidentally' mistaking a kilo of sugar for a kilo of cocaine at a Woolies supermarket - it just isn't something that can possibly happen.

So aren't we so glad that our pretender to a Big Brother government is going to spend \$140M to prevent that from happening.

And I am so looking forward to our chocolate ration increase from 30g to 20g too.

What arrogance they have to think we are such an illiterate society that their machinations are not so totally transparent.

Here is the germane part of my response to the pilot request. My goals being twofold:

1. Point out the waste of the whole exercise
2. Should the whole waste of time exercise go ahead, demonstrate that it can be done (pointless as it is) at minimal cost.

Please excuse the poor scan quality.

I doubt we will ever know for sure, but it would be 'oh so' interesting to see just what tax payer claims for funding other

ISP's put in for the pilot. I guess if we see any radical 'special offers' for broadband pricing in the near future, we will know just where the money really came from for that.

Posted by Steve Waddington at 15:57

Monday, November 10, 2008

HSPA Road Trip - Day 2

The first stop for the day was my friends block at Napier - a small community 25km north of Albany along Chester Pass Road. We could get a 40% signal from the farm gate. Using a standard car mobile antenna on the top of his shed, it is possible to get a useable signal from there as well.

With a Car Antenna on the shed roof, coverage is possible from the Napier farm block.

Signal was lost at 28km north of Albany. We then turned onto the Porongarup road towards Mt Barker and re-acquired a signal from Mt Barker at Dukes Vineyard 15km east of Mt Barker.

Dukes Vineyard is the absolute limit of coverage from Mt Barker. Signal was dropping in and out.

Back on Chester Pass Road, there was no signal through the Stirling Ranges and Gnowangerup until we reached Broomehill.

There was no problem with the signal at Katanning, but was lost another 8km north.

Full Signal as we pass Katanning

Wagin also presented no problem for coverage.

Wagin was worth stopping to look at the beautiful old buildings - Hotel, Court House and Town Hall.

At Narrogin, there was full coverage, and also the speed was back up to broadband speeds of 2Mbps+

Narrogin Exchange Hotel. I really like these old pubs in the wheat belt.

Signal strength was also fine at Brookton. We lost the signal for about 20km at 20km west of Brookton, but then came back into range 91km from Perth. We had full coverage from that point right into Perth along the Brookton Highway.

Brookton Club Hotel. Brookton was our last stop before heading back to Perth.

Summary of Day 2

The total trip distance was just shy of 1,300km. 'Large' towns like Katanning and Narrogin have populations of less than 4,000 people. Brookton, Denmark, and Wagin are more typical of the smaller south west towns, with populations in the 500 - 2,500 range.

Trip route in green and black. Green = Usable Signal, Black = No Signal.

Posted by Steve Waddington at 19:39

HSPA Road Trip - Day 1

Setting out from Perth at just after 6pm we headed south down the Kwinana Freeway to Halls Head, Mandurah. The one concession to testing HSPA as we were traveling was to use the USB cable that ships with the SIM stick, and some velcro tabs to stick it to the dashboard.

The following testing was done continuously for the trip:

1. Active MSN messenger
2. Active Youtube streaming

3. Email checks every 5 minutes
4. speedtest.net tests at each stop and periodically en-route
5. VoIP calls to a mobile number in Perth

For the hour long trip to Mandurah, the signal strength was never less than 90% and download speed was consistently above 2Mbps.

On Saturday morning, only an hour later than I planned, we took Highway 1 to Bunbury. The HSPA signal continued to be above 90%. We stopped in Bunbury for coffee at the Walkabout Cafe. Download speeds were 2.3Mbps.

Bunbury-Busselton-Margaret River continued to be no problem. We didn't drive into Busselton, but stuck to the highway and the bypass. I figured that if we were getting 2Mbps downloads from the highway, it would not be any worse in the centre of town.

At Vasse, still with perfect signal and 2Mbps+ download speed.

Roadworks on the way to Margaret River, at the Metricup turn-off - still 2Mbps+

Predictably, Margaret River presented no problems at all for access or speed.

As we headed towards Caves Road from Margaret River, the signal dropped out after 5km after the Wallcliffe Road turn off. For most of Caves Road there was no signal in the valleys, and a weak 10-25% signal on the rises.

At the Howard Park winery, we could get patchy access at the cellar door - 300kbps download with 5% packet loss.

The story was the same at the Cullen winery, 60% signal strength from the front gate, no signal at the main building.

Some dope had managed to leave the road maps behind in Perth, so I was relying on Navman to get from Margaret River to Bridgetown. Personally I would have taken the Vasse Highway and then the Brockman Highway, but Navman had different ideas, directing us down the Williamson 'Road' towards the South Western Highway.

It was amazing to see the HSPA signal continue strongly for 15km, then drop out as the track dipped between two hills. On the rise before the drop-out we had 45% signal strength and 1.8Mbps download speed.

Williamsons Road - 45% signal strength and 1.8Mbps download speed.

We regained signal 15km north of Bridgetown, and we were back at 2.2Mbps download speed until 8km south of Bridgetown where the signal was lost.

The Bridge, at Bridgetown.

Signal was reacquired 5km north of Manjimup. At this point the data rate was only 30kbps. This was the same for the rest of the Lower Great Southern. It seems that any exchange in the Katanning call catchment area was not able to provide HSPA speeds - instead, even with 100% signal strength, only dial-up modem speeds were available.

Entry to Manjimup

Signal was lost 10km south of Manjimup as continued on the South Western Highway to Walpole. 4km north of Walpole we were back to 70% signal strength.

Walpole Main Street - AKA South Western Highway

We lost signal again at Bow Bridge but were back in service just after the Parry Beach turn off 20km west of Denmark.

Denmark across the Denmark River

From Denmark to Albany signal strength was never below 75% and we were constantly in service.

We arrived in Albany just after 7pm and the end of Day 1. We stayed near Middleton Beach, in the lee of Mt Adelaide. Mobile services for Telstra and Optus now only two bars from this location. The HSPA signal was also only 40-60%, however there was no loss of service.

Albany Town Hall around Sunset

End of Day 1 - At Middleton Beach in the lee of Mt Adelaide

Summary of Day 1

Posted by Steve Waddington at 12:32

Thursday, November 6, 2008

Using HSPA VoIP

Mobile calls for the last 2 weeks from Perth using a Nokia N96 with Exetel HSPA and VoIP

I have used the published VOIP Telephone Charges for the call cost calculation. They are different to the HSPA VoIP costs here, because the VoIP client is set to use our office PABX rather than the VoIP service servers.

Thanks to a comment made about this blog post, I am adding some addition detail that may be useful.

In addition to the VoIP call costs there is a \$25 once off activation charge and a \$5 per month access charge, the plan prices are here. There is also a 1.5c per Mbps cost, which works out at 0.228Mbytes per minute of talk time, which is close enough to 0.35 cents per minute, or roughly 1 cent per 3 minutes.

There is this useful table on the Cisco web site that gives you the accurate data rates for various compression types. We also have a tool in the user facilities area that will calculate the same thing.

Posted by Steve Waddington at 12:01

Wednesday, November 5, 2008

Media Converter and GBIC fail at the same time

It appears that both a fibre to GE media converter and an ESR line card GBIC both failed last night. When only one was replaced with a spare, the circuit stayed down, but when both were replaced the circuit was returned to service.

It is highly unlikely that two electrically isolated components would fail at exactly the same time. Since both are rated at an MTBF of 90,000 hours, the chance of them both failing in any given hour, after less than 10,000 hours of operation, would have to be in the region of one in one billion.

So I think we can count out that that is what actually happened.

It seems far more likely to me that the problem is actually with the line card itself, that caused a degradation in the signal. The degradation was enough to drop the signal below the level the older GBIC and Media Converter could deal with, but not quite bad enough that the brand new replacements could still handle it.

A replacement line card has been shipped from Cisco stocks and is due first thing. Then we will need to schedule a maintenance window to swap out the suspect one. At that time we will be able to test if my theory is correct.

Since the line card, media converter and GBIC combo are all working now, and we have or will soon have the spare line card on hand, there is no need to rush the swap out. If we schedule it early Monday morning, everyone will be rested up, and we can carry all the traffic at that time of day on the other line card - so there wont be any service impact.

It's been a rough couple of days. The Service Controller server for the NE blew up (literally - burning smell from the mother board) on Tuesday night, and William and Jing worked from 7pm to midnight to replace it. Then this last night, where we all worked on it until about 2 this morning.

Now I am really apprehensive - waiting for the third thing....

Posted by Steve Waddington at 10:04

Tuesday, November 4, 2008

Perth POP Upgrades

A new router is going into the Perth POP today. This will terminate a new cross-connect link to Optus, which we will use to supply an extra 50Mbps of bandwidth for WA.

I will need to bring down the WAIX link, that port will be used for a cross-connect to the new router, which will then terminate both WAIX and Optus IP links. Once the WAIX BGP session drops, all traffic will be delivered via the Perth-Sydney link until the new router is brought into service.

The good thing about this change is it is not going to effect any customer traffic, and more importantly, I expect to have it done before race 7 at Flemington. The exact science of race betting means two choices are already made for me - Barbaricus and Profound Beauty, and I think Mad Rush to close the gate on a 3 way quinella.

Posted by Steve Waddington at 12:40

Monday, November 3, 2008

Filtering Inappropriate Content

I am sure people's (in this case 'people' those who run ISP's) fears are well founded regarding the huge cost white list/black list filtering of inappropriate content. But, although the fear may be well founded, the reason for that fear I can only put down to plain stupidity on the part of those people. And I could be completely wrong - maybe my feeble mind just isn't able to grasp the full scope of what such filtering entails.

I am sure every vendor with parasitic interest in playing on general public fears of child pornography; every sysadmin/IT manager that sees their self interest in building their empire of ever increasing big iron hardware; and every vote grasping politician who sees it as a platform to promote themselves as a champion of family values; will agree - I am a clueless idiot who just doesn't understand.

So why don't I just step through it, and you can judge for yourself where my reasoning breaks down.

Look at the proposed blocking method - black and/or white listed sites using DNS blocking, there are two issues, and I will deal with the trivial one first, which is the supposed huge cost.

In a recent Comms Day article, Mr Lindsay estimated the cost of such filtering would run to \$66M. I worked with John Lindsay briefly circa 1997/8, and always thought of him as a sensible, level headed person. So I have no reason to doubt that his figures are anything but accurate, and based on reasonable assumptions and extrapolations.

Yet such a cost for hardware and software alone does not seem reasonable to me. RBL's and DNS RBL's have been around for more than 15 years. Used mostly for spam or the prevention/mitigation of botnet and other DoS attacks, as an operational utility they are well understood (the use of RBL's over time has shown they are in fact useless at preventing anything, but my point is that, useless or not, their operational use is well understood by operators). They impose some, but not a great, processing burden on servers, and in any event, server processing power is now days easily able to accommodate the modest requirements of these programs.

Essentially all DNS servers have the RBL ability inherent in their setup, and most ISP use it in some form or another. Expanding the use to encompass legislatively required block lists might mean upgrading the servers, but I would not think the cost would be more than a few thousand dollars, even for a very large ISP.

So where is all the cost then? I can only think that it is in the licensing fee that the companies providing the RBL service would charge. Or, if they are promoting a proprietary system, in the over provisioned, and over inflated hardware, software and licensing cost bundle that they are looking to charge.

My point here is that no such cost is needed. Every ISP has the ability to adapt their, already there, RBL's to update from a central government site (should such a thing be set up). And that would be at no, or a very trivial, cost to the ISP.

Why would a genuinely concerned government, who really had the best interests of the Australian public, not do that? Why would they instead choose to mandate the use of, off shore, very expensive proprietary services? Have a think about that one.

We see then that a DNS based RBL solution does not need to be expensive - which is the trivial issue in the argument. The non trivial issue is that such an RBL solution actually wont block anything anyway.

Blocking the DNS lookup of a site does not in any way prevent access to that site. I am sorry if I now bore the technically clueful with this, but for the others, a simple explanation may be needed to support my claim.

The first thing we have to ignore is that most inappropriate content does not come from web sites in any event. Rather, it is hidden deep within news groups, send via closed IRC chat rooms, or uses other clandestine methods for distribution. Undercover FBI agents can infiltrate and track such perpetrators, but there is nothing a network operator can do, no matter how much blocking and restrictions are put in place.

But let's put that aside, because it is not germane to this particular argument which is only focusing on blocking the 10%

or so of 'obvious' bad content that web sites represent.

We must also dispel the misconception (assuming there is one) that some innocent web surfer might 'accidentally' stumble upon such sites. You can do your own investigation, or take my word for it; such a thing is just not possible - for obvious logical reasons.

The purpose of such a block is therefore to prevent people who are actively seeking to either distribute or obtain such material.

DNS addresses are simply an alias for the underlying web site IP address. For example, www.exetel.com.au is much easier for humans to use than the IP address of <http://220.233.0.9/>. Yet if you know the correct url string, using the IP address instead of the domain name will give you exactly the same page.

If you block the DNS address, it only prevents access by someone using that human readable name, it does not and can not prevent access by someone using the IP address of the server.

The point here is that NO DNS RBL CAN PREVENT ACCESS DIRECTLY VIA THE IP ADDRESS.

Which means that any scheme to use RBL's based on DNS black/white lists is completely useless, because it can be completely bypassed by anyone prepared to type in 12 numbers instead of a string of letters.

Even though the cost, if implemented by a genuinely concerned government, may be small, it is never the less pointless, because it can be so easily negated.

This is why I make the statement that those moaning of the cost of such a scheme must be stupid. They just seem to have really missed the main argument that it doesn't matter what it costs because whatever the cost, any DNS RBL based scheme will not block anything anyway.

My final point is that while such a DNS RBL scheme has no merit, there is in fact a very simple way such a block can be achieved.

For many years, at Exetel and companies I have worked for prior to that, I have used the expedient and simple 'null 0' route to block a range of different site. Typically these have been servers originating spam, botnet, DoS, hacking or other potentially damaging and unwanted traffic. And sometimes I have even used it as an emergency measure to block on-net IP's sending spam or other damaging traffic out to the Internet. I can tell you that a null 0 route is 100% effective, and there is absolutely no way around it, short of hacking into, and removing the line of code from the router (and there has never been a case of a Cisco router being hacked, AFAIK). Further, the amount of server and network resources it requires is as close to zero as to make no difference.

Automating the implementation of such a null 0 route list is also trivial. Far easier than a DNS RBL. We have a such code, which is about 100 lines of Perl, in place using NetSense that we use for traffic management and automatic code balancing - far more sophisticated than what is needed for a null 0 RBL. A simple modification would be all that is needed. Even starting from scratch I can't see it taking any competent coder more than a few hours to produce the twenty or thirty lines of code needed to do it.

Posted by Steve Waddington at 11:28

Sunday, November 2, 2008

Pocket Money

I had an idea to have a 'mini-agent' class of reseller; somewhere between getting a credit for recommending Exetel, and being a full agent for Exetel. Mainly my idea derived from my kids looking for ways they could get extra pocket money - the novelty of delivering marketing material to letterboxes having worn off.

The opportunity arose recently to test the water for such a scheme, at our local school fete.

With a few pointers in the right direction, my two eldest set themselves up as Exetel Agents, used the Exetel web site to come up with a brochure and other 'marketing material', and set themselves up a stall at the fete as you can see above.

I curtailed the idea they had of using our home inkjet printer to print out the 150 flyers they thought they would need. Instead I took them to Officeworks to run off the copies for 6c each - the major outgoing for the exercise.

From my sons point of view, they thought the project was successful - handing out about 100 flyer's to people who expressed interest. We don't know yet how many will actually sign up, but even if it is only 5, they will have achieved their modest financial objective.

From my point of view, it was the important first step - to see what would be needed to make the scheme generally available. Which is:

1. Create a 'mini-agent' class that gives a once off commission for sign-ups with that code
2. Create a 'kit' including flyer's and brochures that can be downloaded and printed
3. Determine a small incentive for customers to sign up using the MA code
4. Create a web site for MA's to register (a pre-requisite being they are already an Exetel customer)

Posted by Steve Waddington at 09:58

Thursday, October 30, 2008

A clean and safe internet

That's what we want. An Internet with no spam, no porn, that is not show and tell for every malady an perversion of the human psyche. An Internet free of inappropriate content and dissenting opinion. Heralding a world where wrong minded people will disappear quietly, and our benevolent regimes can inconspicuously get on with their final solutions of the day.

I think that satirical paragraph covers both arguments for and against Internet content filtering.

If the law of the country of operation of a network owner required the use of content filtering, then a network operator on their own would be very foolish not to adhere to the requirements of that law. Personally, I might make my protest elsewhere, but I wouldn't jeopardize the welfare of my company, partners and employees, nor the service to our customers, on the basis that my personal views may not match those of the current legislation.

However, there is a point at which, for me at least, I don't need legislation to tell me what is right or wrong. Some things are just plain wrong, and I will not carry traffic for them on my network. I don't care what opinion others have, there is no moral ambiguity in my mind. I am capable of, and will, make that decision on my own, and require no nanny state to be my conscience on that.

Ok, my moral rant aside, just what is all this hoo-har over content filtering anyway?

We received an email from a content filtering company a few days ago. They had clearly gone to the trouble of informing themselves of who the directors of Exetel were, and in their email gave a lengthy spiel about the, possible, upcoming requirement all ISP's in Australia will have to filter content in the future.

Their claim was that they had the best (or should that be bestest?), easiest to manage, easiest to deploy content control system that would filter "'illegal' as well as 'general inappropriate' material". And their email was to find out our interest in running a 'trial' of their system. (Although exactly what we would trial, since AFAIK, there is no authority providing legally defined 'illegal' or 'inappropriate' traffic in Australia. I guess we would just have to trust them on that).

My brief reply was a) we would comply with legislation, when and if there was some and b) we believe we have an easier and lower cost method of content filtering, that does not need proprietary hardware and software, in any event. I asked them to clarify what they meant by 'free trial' and let me know what the cost/licensing fees were for their solution.

I received back today a six paragraph email, again lauding the benefits of the filtering system, with another marketing brochure attached, and an explanation (in the last paragraph) that

"As the costs for xxxxxx and xxxx Technologies are scalable and depend on traffic volumes and differ between implemented and hosted solutions can we arrange a time next Monday or Tuesday for myself and xxxxx xxxxxxxx our MD to call and discuss the technologies and indicative pricing in more detail."

I can tell you, that when someone goes to that length to NOT give you the price of something, you can be 100% sure it is going to be very expensive. My guess would be in the order of \$5-\$10 per month per 'license' (read; per end user). A very, very expensive RBL in other words. And we all know how well RBL's work to prevent unwanted content - that is why there is no email spam any more.

Posted by Steve Waddington at 09:30

Thursday, October 23, 2008

Lawyers, Guns and Money

Fortunately, nothing has hit the fan yet, but if it were to, then we can take some comfort in the preparation we have taken - with lawyers and money at least. I guess the guns will have to wait for the revolution.

What the hell and I talking about? I met with our solicitors yesterday to review the opinion on copyright infringements we have obtained from an SC (which took care of a large portion of the money mentioned above). The 19 page advice was, John and I thought, straightforward and as cogent a piece of reasoning as could be asked for, that anyone should be able to follow. Hopefully even a judge, should it come to that.

It contained references to case precedents, and a little bit of legal jargon, but was easy enough to follow read in context. It said, pretty much, that what Exetel was doing with regards to forwarding copyright infringement notices, then assigning the end user to a captive portal for 2nd and subsequent notices, was, in the SC's opinion, legally, the right way to handle the issue.

An interesting comment that our solicitor made was that when they had first engaged the SC, he commented that he would 'given his opinion, but did not think it would provide much comfort to the client'. However, after seeing how we handled copyright issues, he was 'pleasantly surprised that he would be able to provide very positive advice on the approach the client has taken'.

My opinion is; I think the money was well spent. It provides vindication for the action, and flack, we have taken, and it shows that, at least in this case, the morally right thing is also the legally right thing to do.

Posted by Steve Waddington at 08:46

Monday, October 20, 2008

The Nokia N96 and HSPA

We took delivery of half a dozen Nokia N96's last week, to trial in house with our HSPA service. And I have to say, the 'phone' is a technological marvel. If you can even call it a phone with its Symbian OS 9, 16Gbyte memory, GPS, media player, WLAN, HSPA, quad band and many other applications. But, you can also make a call from it, so 'phone' it is.

I said a fond goodbye to my four year old, completely theft proof, old Nokia, but technology marches on.

Over the weekend I have been comparing two different VoIP clients - Fring and Nimbuzz, with the standard mobile call quality. With six calls from Sydney to Perth, the results were:

GSM MoS: 3.8
Fring MoS: 3.0
Nimbuzz MoS: 3.7

Three calls to Sydney based mobile numbers gave:

GSM MoS: 3.8
Fring MoS: 3.5
Nimbuzz MoS: 3.8

I would have liked to have tried the native SIP client in the N96, but could not get it to work. The Nimbuzz client had one buggy issue - where every second call required the client to be shut down and reloaded. Not sure if that is a software bug with the N96 or my inept setup. In any event, the quality of Nimbuzz was certainly better for calls where the latency budget was tighter.

We are also getting in some Huawei HSPA phones this week, so it will be interesting to see how the two brands compare.

Posted by Steve Waddington at 09:48

Friday, October 17, 2008

A Day off to Watch Cricket

I don't know what it is about those ads that I find so offensive. First we have the rat infested house and now the 'Australia Day' ad.

'But that's funny dad' my youngest son said to me.

Sure, for an eleven year old, I guess it is funny. What is there to find offense with in that?

Two things;

First, like that whole series of ads, is the character of the thicko 'Homer Simpson-esque' father. The creators of the ads haven't made him just stupid, they have given him an IQ of 70 because they think that anything higher will be beyond the level of their target audience to recognize as stupid. In other words, they are saying that they believe the people who will buy the product based on the ad are almost, but not quite, as stupid as the main character.

I don't know, but if I lived in the western suburbs (which I used to) or a country town (which I did), I may well be offended by that. Which is probably why I am.

Second, you swillers of inferior over priced chardonnay, in your BMW'a and private education your parents worked hard to pay for for you. Just what the hell do you think January 26 does mean?

A day off to watch cricket, have a barbeque with your mates, a few beers, or whatever it is you traditionally do on that day is exactly what that day should mean.

We don't 'hail he chief', we don't put our hand on our hearts for the oath of allegiance and fly our flags high. There is no rockets red glare here (ok, maybe at the Perth Sky Show). And it is no longer a concern of ours if God saves either Queen or Governor General.

Australians all, we rejoice with our cricket, tennis and barbeque's.

The assignment question was 'What does January 26th mean to Australians?'. The kid and his father nailed it. But if obviously means something quite different to the wankers who wrote the ad and the Americans who approved it.

Posted by Steve Waddington at 12:20

Thursday, October 16, 2008

Another thing

There is a funny, or crude, or crude and funny, joke where two blokes see something very rude and male chasing something else very rude and female down the street. One bloke turns to the other and says 'It's just one ***** thing after another'.

But isn't it just.

I'm talking about our little 'hiccup' yesterday. A one hour outage for all services.

This isn't an exact quote, but it is essentially what the co-lo data centre manager said to me at the time:

"A contractor was adding a new GPO to UPS circuit 6. The bus bar was a little bent out of shape so he tried to force it, but when he did it shorted against another phase rail. That overloaded the UPS which went to bypass and also tripped circuit breaker of the bypass power supply"

For crying out loud. I mean just get past the string of events that happened to cause it, then consider that there are five other circuits that it could have happened to instead of the one our racks are on.

It's cold comfort, but if someone else doesn't I am going to get that idiot contractor fired. He's a menace to everyone.

But wait a minute.... I just remembered, we pay that data centre quite a bit more for each rack so we can have a second power rail on a different circuit - to safeguard against exactly that sort of event. But both power circuit were out on all three racks.

So now it is obvious they are connected to the same power rail, and same UPS after all, and they have been charging us for a redundant supply for the last two years and not delivering anything but extra GPO's on the same supply. I think a reasonable person would call that a rip-off.

[some quite offensive opinions deleted]

Fan-bloody-tastic. Between idiot contractors we have no control over and rip-off suppliers, what chance does anyone have?

Posted by Steve Waddington at 08:34

Monday, October 13, 2008

Why HSPA beats wifi

I have been a big fan of wifi since I first used it at an Interop convention in Atlanta in 1997, and then marveled at how useful it was while fixing a network problem in Australia while using the free wifi in Singapore Airport en route to Taipei in 1999.

Surely access to cheap, or free, ubiquitous Internet anywhere in the world would only be one or two years away at most.

But it never really happened, did it. Despite every laptop post 2002 having built in wifi, and, at least in the US, wifi being available everywhere there was a coffee shop, a McDonalds or a progressive municipal government.

In the home it works just fine. But even for a small office LAN, we had to abandon wifi as a viable medium when we were just an office of six - because everyone else in our building had the same idea and contention for spectrum with the 12 other networks in reach of ours meant drop outs between every hour and every few minutes.

There may be other, or many other, reasons why wifi never fulfilled its promise. But for me the main thing was, while convenient in most places, it was just that little bit too inconvenient and just that little bit too unreliable for it to be relied upon.

And while it may be argued that wifi is 'free', the fact is it rarely is.

Although HSPA certainly isn't free, it is in fact cheaper, far more reliable, and far, far more convenient.

Another big plus for HSPA is, at least for Exetel, the IP address is fixed. Very easy for firewalls and security policies to allow. Compared to some random third party network needing many different layers of security and encryption to be maintained if external access is a requirement.

Look at a wifi data path:

your laptop -> local wifi 'coffee shop' (with whatever restrictive NAT and firewall rules they use) -> coffee shop service provider -> service provider's Internet supplier -> your ISP's Internet supplier -> your ISP -> your company's Internet access -> your company LAN/firewall

vs the HSPA data path:

Your laptop, with one fixed IP address -> Exetel -> Your company LAN/firewall

If you have (as Exetel does), a SIP capable PABX, then a combo headset and VoIP client on your laptop also gives you your office dialtone anywhere you are in Australia (that a mobile network reaches, obviously). Or if you want to get really sophisticated (read; get your teenage son/daughter to set up your HSPA mobile), you can get office dialtone on your mobile phone.

Which means exactly what you think it means. No extra cost for calls from mobile, just the same rates as if you made a call from your office phone - because effectively, you are.

Or, if you don't have a PABX (or an IT guy) that can do that, use a VoIP provider (Exetel being the obvious choice for the Exetel network :)).

But wifi is FREE! Sure. Unless your staying in a hotel, where it is \$20 a night. Or trying to find a signal in the Rundle Mall in Adelaide, where it is \$6 per hour through a local wifi provider. Or in Albany or Broome and there is no signal anywhere.

National portability is one thing, now the next thing I need to investigate is global roaming for HSPA, which I doubt will be cheap. I noticed however that Sri Lanka Telecom had international roaming for HSPA, and I see Optus is part of the Bridge Alliance. So I think the cost will be relative. It will be interesting to see how it compares to the \$20 per 24 hours hotels seem to charge everywhere in the world.

Posted by Steve Waddington at 11:09

Sunday, October 12. 2008

HSPA VoIP and Latency Effects

Over the last three days I have been testing VoIP over HSPA, alternating with connecting the family 'games' computer via HSPA to test selected games performance.

For VoIP I have been calling mostly locations in WA and mobile numbers. Since the VoIP switch is in Sydney, it means the call 'trombones' from Perth to the Sydney switch, then back to Perth again. I selected this route because it is the worst possible path, latency wise, for a VoIP call in Australia.

Using the ping utility from my computer to Sydney, the latency is recorded at 160-180ms via HSPA - higher by 80ms on average than my ADSL LAN connection. I have commented in the past that the latency reported by ping seems out of kilter with the actual data performance. It is always better to start from a lower count, but even so 180ms isn't going to blow the latency budget for any generally available codec.

Gaming however was going to be the real acid test. 80ms extra latency is a life or death matter, as I was told, among other protests, such as how uncool it would be to let 'mates' down with a laggy connection if the very important (and seemingly endless) Team Fortress II game.

Tests were as follows: Friday & Saturday 10am and 2pm VoIP calls to a Mandurah land line, Albany mobile and Perth mobile. Friday 3pm to 5pm and Saturday 10am to 11:30am on-line games.

Also, as I type this at 10am Sunday morning, I am in Adelaide airport waiting for my connecting flight to Sydney, and I have just completed a call to my home in Perth via VoIP, which was also fine (MoS 3.8).

The results for VoIP were pretty much as I expected, which was a MoS score between 3 and 4 for all calls. One call to a Perth mobile dropped out after a few seconds, but was completed successfully on the redial. I could have been a mobile phone drop out, or it could have been VoIP, but as it was a one off event, it was not a concern.

So I think I can safely say that that places VoIP over HSPA solidly in the 'usable' range.

The game test however went better than I thought it would. TF2 seems to use both UDP and TCP ports, so while I was sure the UDP parts wouldn't be a problem, I wasn't sure that whatever part of the game TCP was responsible for wouldn't cause a (fatal) problem.

But it didn't, or if it did, not enough to be noticed. Richard reported that there was no real difference as far as he could see. So too with the other games the younger two play such as Maple Story, Gunbound and Guild Wars, all reported no noticeable difference.

One puzzling thing however is that all testers, not just me, have reported 'much better' performance since the migrations on Thursday night from Internet access via Optus to via Exetel. I don't know how that could be on a common shared backhaul network. It suggests that the backhaul bandwidth may be partitioned so that Exetel, with as yet very few users, have free range of the allocated bandwidth (if it is partitioned in that way).

It is something I will look into further next week. It doesn't make any difference at this state how it is done at the moment, and shouldn't make any difference in the long run - never the less, one of those things that needs to be nailed down.

Posted by Steve Waddington at 12:28

Friday, October 10, 2008

Where have all the [Tech company sector]'s gone..

.. long time passing? (Apologies to Pete Seeger and Joe Hickerson)

I 'grew up', career wise, in the 80's at the start of the PC boom. Ah yes, the days of the NEC APC III, with an 8Mhz CPU, twice as fast as the IBM XT, for the same price of \$6,500.

As the market for PC's developed through the 80's, companies like Computerland, HiSoft, Imagineering and Ferntree grew into very large businesses. With hundreds of staff and hundreds of million dollars of turnover.

Their rise what the beginning of the end for a previous 'darling' of the technology sector - the mini computer. Few people would not remember DEC (before Compaq and before HP). But who recalls MIA, Prime, Pyramid or Wang? I remember walking through West Perth circa 1985 and marveling at all the shiny signs and prestigious offices of those mega-companies of their day. Usually headed by sensible Harvard graduates and with rock solid trading credentials, their future was assured in an increasingly technology hungry world.

Just like their successors, the large corporate PC companies, it seemed they would be around for ever.

Why this boring reminiscing?

I think it is because, with recent stock market events, there are enough parallels to trigger in my brain an association with the large computer companies of yore and the large ISP's of today. I can't put my finger on all of them, but the obvious ones are that they have a headcount of hundreds or many hundreds of staff and having a major supplier that is also a competitor.

In a tightening economy, a supplier/competitor looks to a) protect its own retail revenue, b) reduce the risk from very large accounts (which include wholesale customers) with more rigorous enforcement of contract terms. So there is a squeeze from both ends as wholesale supply is tightened and retail sales come under even more competitive pressure. Even the large supplier/competitors have found it necessary to retrench new staff, lay off contractors and can new projects as they in turn are affected by the price competition they have, in the main created (through the undoubtedly unethical but questionably legal practice of subsidizing the retail operation so it can sell at a price less than the wholesaler can buy for).

John has been through at least three iterations of that cycle (Mainframes, mini computers and PC's), and I have been through two, and it has never been very pleasant. Obviously it isn't the same in each case, and the differences far outweigh the similarities. But I think we have both had that cycle at the back of our minds as we have developed Exetel.

Certainly a lot of our planning and the systems we have put in place, whether purely coincidental or not, give us the low operational cost base and flexibility needed in these times. I think I would rather be us than our competitors at this point in time.

Posted by Steve Waddington at 11:28

HSPA Looking good

The Exetel HSPA service was swapped over from the Optus Network to Exetel's network last night. The change over was pretty smooth - just had to change one setting from the default to 'exetel1' in the HSPA setup software.

Download speed are around 2.8Mbps average - faster than when it was tested on the Optus network from my home location - but it could be a time of day difference.. Download from the Exetel servers is a little quicker, as is access to the Exetel Intranet at the Sydney office. Which is to be expected.

Pings are 165-170ms back to the Exetel servers (from Perth)

Nice to see this service located on our own network at last. I am going to give it a workout with VoIP in the morning and see how that goes. Then I will turn it over to my children and find out how it stands up to latency sensitive on-line games.

Posted by Steve Waddington at 00:27

Tuesday, October 7, 2008

Cache Boost

The changes on Monday had about the effect I expected on cache generated traffic. We measured an increase of 100Mbps average.

Also evident, and as expected, is the far lesser impact that the NetEnforcer time of day rules have. As other planned capacity upgrades come into play, I am anticipating seeing another 150Mbps increase to bring us up to 400Mbps average traffic generation.

Posted by Steve Waddington at 17:51

Monday, October 6, 2008

HSPA Road Trip

I had planned it for this weekend, unfortunately Internet access on the Exetel HSPA service is still via Optus, and it won't be switched to our own network until Tuesday night. I still could have gone, since the Internet access part will make no difference to signal acquisition, however I want to be able to report exactly what an Exetel user will be able to experience from the service. Since I will be in Sydney for the next two weeks, the trip is postponed until 1 November.

So here is the travel plan:

Day 1

Stage 1: Leave Perth early Saturday morning and drive down the Kwinnana Freeway to Manjurah.

Stage 1a: Morning tea at my Mom's house - which has patchy mobile phone access - usually 2-3 bars.

Stage 2: Mandurah to Pinjarra and then down the SouthWest Highway via Harvey to Bunbury and on to Busselton

Stage 2a: Lunch in Busselton

Stage 3: The Bussell Highway to Margaret River and selected wineries

Stage 4: Brockman Highway to Bridgetown via Nannup. I want to test access in Nannup - last time I was there it was the black hole of all mobile phones, had to make a call from the public phone box in the main street (and was hard to hear over the banjos).

Stage 4a: Coffee break in Bridgetown, in the heart of the Central Ranges on the banks of the beautiful Blackwood river, one of my favorite spots in the Southwest.

Stage 5: South Western Highway (Highway 1) to Manjimup (resisting the urge to take Muirs Highway to Mt Barker - gone are the days when that was a one hour drive) through the Shannon National Park and on to Walpole and Denmark

Stage 5a: Dinner in Denmark

Stage 6: Short drive to Albany. End of Day 1

Day 2

Stage 1: Albany Highway to Mt Barker, selected wineries, Porongurup National Park then Chester Pass road through the Stirling Ranges to Gnowangerup

Stage 1a: Lunch in Gnowangerup

Stage 2: Great Southern Highway to Katanning, Narrogin and Brookton

Stage 2a: Coffee in Brookton

Stage 3: Brookton Highway to Perth. End of Day 2

For each stage I, or rather my traveling companion, will be testing:

- web browsing
- download speed
- access to email
- VoIP quality

I am very interested to see what the service quality is in the moving vehicle. Also how the coverage will go in some of the

hillier and remoter areas of the South West. I fully expect the signal to drop out in some of those places, around Nannup and from out of Manjimup through to Walpole. What actually happens will determine just how effective HSPA can be for overcoming the sparseness and distance tyranny of datacoms in Australia

It would be a dreadful trip in a 30 year old Combi, which I am sure would be just fine for the epic journey from Woolongong to Sydney, but a real Australian road trip? - yeah, sure. My regret is that for this 'I've been everywhere' 1,500km trip, I no longer have my old ZF Fairlane. I guess the Honda will have to do.

Posted by Steve Waddington at 12:11

Sunday, October 5, 2008

Bandwidth Use Jump

We added another 200Mbps of bandwidth in Sydney, as well as 50Mbps in both Brisbane and Melbourne last week. To see it instantly 'sucked up' into a new peak bandwidth high of 2.9Gbps the day it was switched on. Which, on those 24 hour figures brings the user bandwidth average peak usage to 52kbps.

Being the end/beginning of the month was most likely the reason for the sudden jump. The last night of the month some people seem desperate to use up whatever remains of their quota. And on the first of the month, people who had run out of quota must breath a sigh of relief and get back to their habit.

Of the 2.9Gbps peak, only 2.2Gbps was delivered through Internet bandwidth suppliers. The break up of the source of the other 700Mbps is:

- 200Mbps from the PeerApp cache
- 220Mbps from Pipe Peering in NSW
- 180Mbps from the Akamai servers
- 100Mbps from other peering (Vic & Qld Pipe, WA WAIX)

Starting this week, we will be reconfiguring the topology of bandwidth into Sydney, with the objective of allowing the PeerApp cache to generate more bandwidth. Now that we have completed the change over to the 6500, we expect to be able to get another 300-400Mbps generated from the cache.

In conjunction with that we will also be trailing the DPI product from Cisco during the month. Putting direct bandwidth into Melbourne and Brisbane, and soon Perth means we will need something at those sites as well. Although, I do need to run the cost model again to see where the break even point is with the promise of decreasing bandwidth costs.

Posted by Steve Waddington at 18:22

Thursday, October 2, 2008

Last to Know

Occasionally, not often, but occasionally, I will get up from the keyboard and do something not even remotely related to Exetel. The reason it is so occasional is that almost always, without fail, something will happen that makes me wish that I hadn't. Hours, days or months can go by in the meantime with no problem at all. But move away from line of site of the computer, and out of earshot of the phone and WHAM, major outage.

So yesterday evening I was out in the front yard, putting in some new lawn seed, foolishly having left my phone on the kitchen bench. I was in the process of admiring how flat I had managed to get the top dressing I had been spreading, when my eldest son came out with my mobile and said 'Phone for you Dad, its been ringing a lot'.

Five missed calls. Gee, thanks son. Sorry about all that ringing, it must have been distracting from your Team Fortress game.

This particular call was from John to tell me that our main NSW Internet bandwidth had just come back on line and did I know what had caused it to be down?

"[supplier] was down?" I asked, or some comment of equal intelligence

Sure enough, a one hour outage from our major NSW bandwidth supplier. Very un-pretty:

My first paralyzing thought was that some fault had developed with the 6500. Chilling, because if that was the case, and unless it could be identified and fixed real quick, it may mean having to roll back to the 7300 routers it had replaced - a daunting task.

There was no problem evident with the 6500, a big relief.

The cause was revealed by the router logs, which showed:

```
Oct 1 18:15:24: %BGP-3-NOTIFICATION: received from neighbor xxx.xxx.xxx.xxx 7 6/6 (cease) 0 bytes
```

Which means the neighbor explicitly reset the session. Subcode six being "Other Configuration Change".

Clearly the supplier had made some maintenance change, probably not thinking it would affect anything. 6pm Sydney time being what, 4am on the US east coast. A usual time for service non-impacting maintenance.

We have asked for a full incident report, so it will be interesting to see exactly what was the cause, and why it took an hour to restore. But of vital importance is to find out how such a thing could happen and not be immediately corrected - vital in fact to continuing to use that supplier as the major source of bandwidth in NSW.

Posted by Steve Waddington at 10:26

Monday, September 29, 2008

Bandwidth vs Subscriber Growth

Having a bit of a 'shuckle' (the thing you do when you chuckle and sigh at the same time) reading this article over my morning cup of tea. The chuckle came from the dichotomy of what must be either ignorance, expressed by those quotes in the article, of how ISP's operate in the US, or their callous manipulation of half facts for their own self serving purposes. The sigh was because neither the journalist writing the article, or presumably the people he has happy to quote, noticed the logic break between subsequent statements.

I couldn't figure out what any of the proponents were really trying to say. I mean really, these are meant to be the 'captains of industry' in this industry. But looking at the picture of those three, it did remind me of something. Eerily uncanny resemblance if you ask me.

Anyway, it prompted me to review our growth figures for the last 12 months for subscribers and bandwidth. I took the daily average of subscriber connections (that is; not total customers, but the customers connected to the service each day), and daily average of bandwidth.

The results were:

Connected customers has increased by 20% since 1 Jan 2008

Bandwidth use has increased by 34% since 1 Jan 2008

I am not sure I have the math exactly right, but it is going to be something close to this - take the 20% connections growth away from the 34% bandwidth growth, and we get a 14% average increase in bandwidth use by all users over the last 12 months.

Looking even further back, in our first year of operation we needed to provision bandwidth at 23-24kbps per user to meet peak demand (and I recall some carrier suppliers at the time commenting on how high that was, since they were provisioning at 10kbps). Now, 4 3/4 years later, we need 46kbps per user.

It looks like, based on offers currently in the pipe line, that over the next year, the price of international Internet access will make the doubled usage the same cost (er, I mean, the price will have halved). In the mean time, domestic backhaul prices have hardly changed (at all since 1985 in some cases). So while international bandwidth cost decreases, the cost of domestic backhaul as effectively doubled (obviously, because use has doubled).

I was going to make the comment that we may be in the situation soon where international bandwidth is cheaper than domestic bandwidth. But then I remembered when a 64k nailed up ISDN circuit was \$24,000pa and international bandwidth was only \$27,000 per Mbps.

Which, when you work it all out, pretty much fits with the way pricing of ADSL has gone, and where it is likely to go in the future.

Posted by Steve Waddington at 10:51

Monday, September 22, 2008

6500 Happy Snaps

Posted by Steve Waddington at 16:07

Sunday, September 21, 2008

Melbourne Upgrade and email outage

Well, the good news is, the Melbourne POP upgrade went without a hitch. The new LNS router is humming away, and both routers are running at less than 1/2 the CPU load before the upgrade. So not much to say there really, except to note everything went according to plan.

If there is one thing the network gods share with their predecessors, the Olympian gods, it is to punish hubris. In fact, I think the network gods go a step further and start their retributive smiting before it ever gets a look in. And they do it with a level of capriciousness that would make Father Zeus himself cringe.

How do you create a very large and fast storage volume? - you use RAID 5 of course.

How do you protect such an array from disk failure? - by using RAID 1, as everyone knows

And so, a combination of those, RAID 5+1, gives good performance for very large volumes, with full redundancy. "The RAID level for the truly paranoid" as it is sometimes called.

But not paranoid enough it seems.

About 8pm on Saturday night email started to queue on the exemail.com.au smtp servers, and at the same time the imap survive hung, due to loss of access to the mail store.

At first look it seemed that two of the harddrives on each of the RAID 1 mirrors had failed at the same time. The RAID controllers couldn't figure out what to do about it, so took both volumes off line.

Well, that's bad, but not fatal. Spares are on hand, arrays can be rebuilt. We have maintenance with Dell for replacement parts.

Andrew made the appropriate calls, got himself to the POP and started the recovery process.

Only to find lots of strange errors being thrown out by the RAID controllers.

To cut a long story short, it turned out to be a firmware bug in the RAID controllers. Some power spike, possibly caused by a failing drive, had caused them to flag disks on both arrays as faulty, and then give up the ghost completely.

Once that was figured out, a firmware patch was provided and the data store brought back into production. Server loads hit 600 for the next 30 minutes as the email backlog was processed, and then all was well again.

There you go, we were punished for our hubris. How is that hubris you ask?

We made the mistake of thinking that the full redundant design, was, in fact, fully redundant.

We had confidence. Confidence is pride. Pride is hubris. The gods punished us.

What were the chances of the original design failing? Well now we can say it was 100% So that's not much help.

Could it have been built to be even more redundant? Yes, of course. The thing is, there is really no commercial way to justify the exorbitant expenditure needed to do that.

What we are talking about is full duplication of a duplicated redundant system. Plus the network switches, plus the L7 switching, plus plus plus, to make it all work. Think not double, but triple the cost.

But maybe that is what is needed. Not because of any sensible commercial reason, rather, in the way of sacrifice to the network gods. No oxen being handy, and with no drainage for the blood in the POP anyway, maybe excessive expenditure is the modern sacrifice needed to show our humility and placate their wrath.

I wonder how I will phrase that in the board proposal?

Posted by Steve Waddington at 14:00

Wednesday, September 17, 2008

6500 Migration complete

The final of the re-planned changes to migrate to the 6500 router were completed this morning.

Two unexpected issues were encountered:

1. All LNS's other than the one in the same POP as the 6500 could no longer reach the radius servers. This was because the radius servers are accessed on 10.0.0.x addresses, and have no default route (a very simple, and very strong, security measure - by only allowing access to specific networks, it is impossible for anyone on a non-routed network to gain access). When the 6500 assumed the core routing function, it broke the path to the 10.0.0.x address space the radius servers are on.

That was an easy one, only took a few minutes to fix once we realized there was a problem

2. Some servers could suddenly not be reached from some locations. The route tables looked fine, and they could be reached from other locations with exactly the same route path. At the same time, other servers on the same network had no such problem. After mucking around with host routes to try and get a good test case, Andrew spotted that the ARP cache for the the affected servers was stale. A quick flush of the ARP tables and all was good again.

Everything is looking good now. The 6500 throughput is 1.8Gbps and CPU is hovering around the 9% mark.

Posted by Steve Waddington at 10:04

Tuesday, September 16. 2008

HSPA Full Test Results

The preliminary results proved to be a pretty good indication of what the rest of the two weeks of testing turned up.

I can summarize the full results (for me) as follows:

- TCP traffic is clearly heavily manipulated on the carrier network by both time of day and bandwidth
- RTP traffic throughput is much better than any ping or traceroute test would otherwise indicate
- I did not find a location/time when VoIP wasn't usable - the MOS was consistently in the 'usable' band anywhere or any time I used it
- Despite bandwidth QoS variance, practical use for web browsing, email or any other normal application was just fine in terms of 'user experience'
- 1Mbps download was average, and speeds up to 1.9Mbps was common

As a user of the service, today I have canceled my Sydney ADSL2 (which I use 10 days a month when I am in Sydney) and Albany ADSL1 (which I use once or twice a year when the family goes to Albany on holidays) service - because HSPA is going to give me the same usable access, but with far greater flexibility.

Posted by Steve Waddington at 12:12

Thursday, September 11. 2008

A better result

The router change went without incident. Traffic dropped a little for all of maybe 5 minutes, necessitated by the BGP swap from one router to the other.

Other than that, everything is looking good. Traffic on the 6500 is now up to 600Mbps, and CPU load is 4% - much more in line with what is expected.

At that rate we should reach about 16% when it replaces the other two border routers, under full load.

6500 Traffic

6500 CPU

On a side note; I see the LHC fired up yesterday, and amazingly enough the world didn't end. Still going to take 6 months of calibration before it is run at full power though.

Posted by Steve Waddington at 09:12

Wednesday, September 10, 2008

Router Upgrade #2

After going back to the drawing board, and waiting a fair amount of time for some new cross-connect cables, we are now ready to make the changes to bring the 6500 into production.

The work is as straightforward as we can make it. The 6500 is racked and powered and connected to the core network via 6 new cables (no swapping cables or routers this time).

The first step will be to replace the function of one border router with the 6500. 24 hours after that, the second will follow, and then if all is still well, the third.

Posted by Steve Waddington at 09:49

Sunday, August 31. 2008

&%*@#\$_!

Great. No sooner had I hit return on the last post, when 5 sms threshold alarms arrived. BGP is flapping on all border routers, caused by `_something_` - no one knows yet.

Latency to the Sydney servers from where I am is 1.2 seconds with 40% packet loss. Totally unusable to access any of the routers and investigate or assist directly with working out what it going on.

I can get email after a fashion, so I can see the messages from customers hitting the forum, but the connection is too useless at the moment to even access the forum. Not even sure if this blog post will make it.

Ah, here we go, pings are back so 120ms. Whew. Now I just need to wait for Vinna's incident report email.

What a horrible start to the week.

Posted by Steve Waddington at 09:20

Early Sunday in Singapore

On my way to Colombo to work with James there to get our comms and other bits and pieces sorted out.

Reading on the plane last night I saw this old quote from Arthur C Clarke "Any sufficiently advanced technology is indistinguishable from magic." (Clarke's Third Law of Prediction).

At the time for some reason I was also thinking of the Fermi Paradox, which Clarke alluded to and Baxter stated outright, isn't a paradox at all, rather, the evidence is all around but we don't know what we are looking at.

Which led me to the thought that perhaps any optimally advanced technology is indistinguishable to nature?

Maybe it could be stated as the Clarke-Baxter nth law:

"Any optimally advanced technology is indistinguishable from nature"

So why aren't these optimally advanced beings enlightening us so we don't have to continue to wallow in the mire of our own human stupidity? Ah, well, you see, they are. The plain and simple message is bombarding us all the time, across every spectrum. But of course:

"Any optimally encoded signal is indistinguishable from noise"

We have the answer then, all we need is a way to decode it. Tricky. I'll have to think about it. The program will take a little while to run. Maybe Douglas Adams knew, but he's not telling.

Posted by Steve Waddington at 06:35

Saturday, August 30, 2008

Infinite, Free Bandwidth at an end

An interesting article here in PC Magazine. In the past I have termed the US the 'land of infinite, free, bandwidth', a sort of semi-joke referring to the far far greater backbone capacity, and far, far cheaper cost (for non tier 1 carriers) of access in the US.

Last time I looked (or rather, noticed an email about it on NANOG), bandwidth was being offered for \$10 per mbps to a small regional ISP who only wanted 50Mbps. One can only assume that larger quantities would attract much lower pricing.

But let's stick with \$10 per Mbps for arguments sake. Even though the price of bandwidth in Australia continues its lifetime trend south, even at the best rates I have heard being offered today, we are still in the order of 15 times that cost.

For Comcast to limit users to 250Gbytes per month must be indicative of some stress, either on the bandwidth or routing capacity of the network, and their engineers provisioning forecast against plotted growth of usage increase. Of course, I could have it completely wrong, but it has to be something like that.

So in today's Australian cost terms, it is like an ISP constraining downloads to 16Gbytes per month. Which, when you think about it, is around the median of broadband plans on offer anyway.

I sort of hoped when I started writing this, that some profound revelation would 'pop out' by the time I got to the closing paragraph. But it hasn't. All I can think is that end user bandwidth use in the US has finally caught up with the infrastructure build of the dot com era, and now we can see where the investment is coming from to fund the next quantum leap in capacity.

Posted by Steve Waddington at 08:13

Friday, August 29, 2008

HSPA Preliminary Results

We have been testing HSPA access in preparation for going live with our own HSPA product. From the initial testing I have done, the results indicate to me:

- Signal strength is no indication of throughput
- there appears to be a good deal of traffic manipulation going on within the carrier network
- Ping and traceroute results do not correlate strongly with download performance
- VoIP is usable even when ping return times would indicate otherwise

Here are some examples:

- From my home in a southern suburb in Perth, I was getting 95% signal strength. Over the period of the morning, ping times went from 350ms to 800ms back to an Exetel server. Download speeds at that time were around 1.9Mbps. In the evening, ping times were as low as 180ms and averaged at 300ms, however the download speed was around 800kbps.
- From a cafe in Perth city, I had signal strength of over 90%, ping average was 390ms and download speed was 800kbps. However from Clancy's Fish Pub in Fremantle (being an HSPA tester is hard work), pings averaged 420ms. signal strength was 45% but download speed was 1.5Mbps.

In all cases there was the odd ping result of 1,000ms +, and an occasional packet drop. Despite that, VoIP from any location was always about the same, rating a MOS between 2.2 and 2.6 each time. (I was using a cheap gaming headset/microphone combo with X-Lite from a laptop. I suspect a better quality setup would easily add .5 to the MOS score). And that was using the worst possible routing path I could think of - calling from Perth to a mobile number in Perth via our Sydney VoIP switch. In other words, VoIP is completely usable, even without a decent headset and codec.

A typical traceroute would show almost all of the latency occurring on the first hop, and quite often the latency improved through subsequent hops until it left the carrier network. For example, the first hop from my laptop might show 420ms, hop 2 470, then 398, 418, 416 etc and then 450 on the final hop. Considering the first hop is, presumably, to a local node in WA, and the carrier network termination is in Melbourne, I surmise that most of the latency is induced due to traffic policing within the carrier network.

Which is fine for TCP and ICMP, which are not latency critical, at least for most business applications. Gaming I would expect would suffer though.

On the other hand, UDP and RTP's seem to slip through the policing cordon, because the VoIP results were essentially the same as my 'control' which was using the same laptop and headset but from my 8Mbps ADSL1 line at home - where my ping is a consistent 90ms to the Sydney SIP server.

We have 10 test units out in the field, and my test alone is not enough to be a representative sample. We should be getting all the results back in the next week, so it will be interesting to see if my preliminary observations hold up.

Posted by Steve Waddington at 16:26

Wednesday, August 20, 2008

Please Explain

I can't remember when I learned this basic rule. I do know it was a mistake I made for at least the first 15 years of my working life.

Simply; any response or explanation that requires more than a [paragraph, sentence, word], is of proportional diminishing value. or perhaps to phrase more formally the value of an explanation is inversely proportional to its length.

So, when you find yourself typing paragraphs of reasons to some relatively junior employee, to try to explain to them why their, to them, very important 'recommendation'; is, in fact, of no value at all. Stop.

It can never be anything but very hurtful for those fragile proto-ego's I am sure. There is just no way, at least for someone of my moderate literacy skills, can phrase any apparent rebuke in a way that will be well received.

Going back to Management Principles 151, saying nothing can never be cause for contention. The corollary of which must be, saying less must always be better than saying more.

I think it can only come down to two views:

1. The people you employ are stupid, and need micro guidance on every aspect of their job (put aside for a moment what that says about you, as their hirer and manager)

or

2. The people you employ are intelligent and capable of figuring out what needs to be done to do the job their are hired to do

For the first case, no amount of verbiage is ever going to help. As the shelf upon shelf of dusty procedures and QA manuals of any bureaucracy will testify.

And for exactly the opposite reason, exactly the same is true in the second case.

Posted by Steve Waddington at 09:07

Monday, August 18, 2008

What went wrong???

I mean, what went wrong...(statement leading to answer, rather than a question).

When things that should go perfectly smoothly go wrong, I can't help but be reminded of my high school science teacher, Mr Savil (a very competent teacher and electrical engineer in his own right), yet who's every experiment in front of the class seemed to defy the laws of physics, thermodynamics and chemistry, and somehow go wrong (often filling the classroom with acrid smoke causing evacuation and the rest of the period off).

Anyway, no smoke came out of the 6500. The physical changes were done by 5:15am, and then it was just a matter of making a few route path changes. Which were in turn completed by 6am.

Everything was looking good route-path wise, and the traffic through the 6500 was matching expected levels. However by 6:30am we could see these problems:

1. Traffic into the network was only 700Mbps, where it should have been 1.6Gbps
2. There was odd packet loss between the 6500 and the ESR LNS's, and much higher than expected latency to the second Sydney POP - 200ms where it should have been 12ms
3. A check of the 6500 CPU load showed it fluctuation from 8% (about what was expected) to 95% over a period of 15 seconds.

The plan was to commence roll-back at 6:30 as things hadn't worked out. The problem was though, that as far as the change plan had gone, and the possible issues we thought we would encounter, everything was spot on. The odd issues we saw (we thought) were probably due to some simple cause, such as a mis-balanced data path causing an interface overload, that could be quickly fixed.

So we delayed the rollback and went looking for the possible error. We actually found two, and with a sigh of relief, made the load balancing right and voila now everything works...

not.

Whatever the problem was, it was not anything obvious. It was 7:20am, and while the deployment team were telling each other what problems we could see, none of us had an answer to how to fix them.

I gave it a few more minutes, but at 7:25 had to order the roll-back.

Very frustrating.

The roll-back took us a little over our maintenance window, and was complete by 8:20am. Although later we were to find some other routing issues that has caused loss of access for some netblocks to some servers.

We broke for coffee at 8:30 and with the main time pressure off, we discussed what, possibly, could have cause the problems we saw. We considered a number of speculative possibilities, but our main consensus came down to some sort of IOS issue. Looking at the traffic graphs between the time we had the routing 'good' and the roll-back, it was apparent that the sum of all traffic inbound to the 6500, which was across 4 GE circuits, did not exceed 960Mbps - where it should have been around 1.9Gbps. Those graphs were suspiciously all flat lining at just under 300Mbps. Which seemed to indicate that there was a 1Gbps circuit bottleneck somewhere in the traffic path. Going through the connections and traffic paths step by step, it is still just not apparent what could cause such a thing.

This one will have to go to Cisco TAC for some deeper diagnosis. And consequently, there was no point continuing the work on the planned Sunday maintenance window, because without knowing the cause, we were bound to strike the same problem again.

A disappointing result.

Posted by Steve Waddington at 09:05

Tuesday, August 12. 2008

There shall be wailing and gnashing of teeth

The biggest single change we have ever done on the Exetel network is in its final stage of planning now. We will be deploying our new 6500 router to replace three border and two core routers in our Sydney POP.

The change itself is fairly straightforward; the interfaces and functions of several routers are easy enough to duplicate on the 6500 router, and we can stage the change over - circuit by circuit - so we can see early if anything goes wrong.

We took delivery of the 6500 three weeks ago, and since then Vinna has pretty much exhausted all the testing and preparation combinations we can think of. Vinna's deployment plan looks sound to me, covering what should be a straightforward change and the likely possible problems that may eventuate. And independently I have covered all the more wild and implausible possibilities I can think of.

Yet, one must never become complacent. It is perhaps an instinct I have developed over the years to let a feeling of trepidation build up that something may have been missed, to further spur my imagination to come up with just what that might be. It makes for some sleepless nights, I can tell you. But if you have ever had that terrible feeling when a catastrophe strikes and you don't know how to deal with it, then you will understand why a healthy dose of paranoia before the event can really pay off.

This particular 'operation' will take place over two days (two early mornings rather). First, the 6500 will be sat next to the existing routers, and one by one, the cables swapped from the 7xxx interfaces to the 6500 ports. That part should take about two hours, so starting at 4am, we should be done by 6. However, the best laid plans etc, so the window is from 4am to 8am, just in case. Should the worst happen and nothing works, roll back will mean just moving the cables back to the old routers, which are still in place and 'hot'.

At 6:30, if everything is ok, checking will be complete and the deployment team can hand over to the monitoring team and stand down. The monitoring team's job is to keep an eye on everything for the next 21 hours. If all is still well, the deployment team will return at 4am Sunday to:

- power down the now disconnected 7xxx routers
- remove them from the rack
- install the 6500 router in the rack
- reconnect and double check its operation

Which will complete the work. Again, 2 hours should be needed, but the window will be for 4 hours, just in case. Then we can all go and have breakfast (nearby). Check again, and we should be ok to let the normal monitoring systems take over after that.

In a way, it will be like a heart-lung replacement, but with one big difference. As bad as it would be, the cardiothoracic surgeon can always lose a patient. In our case, we have no such option; our patient CAN NOT die, or indeed, the time of great tribulation will be upon us, not to mention eternal damnation.

I think we are ready. A bit more checking and then we post the maintenance notice.

Yea, though I walk through the valley....

Posted by Steve Waddington at 12:38

Monday, August 4, 2008

Copyright Theft - The Accused and The Accuser

I read this article with, frankly, not much interest. Same old tired rhetoric.

"Your users stole our copyright material" blusters AFACT

"We aren't the police" Whines the ISP "An besides, it's impossible to do anything about anyway"

Well, the first part of the ISP's statement is certainly true. Although the second part is either made by someone ignorant of what is technically possible, or someone hoping their fib will be believed by someone who is technically ignorant.

As far as not being the police goes, and no one seems to have spelled this out so far, the ISP's dilemma is this. 'Someone' claims that a customer of the ISP is stealing something of theirs, and they request that the ISP prevent the customer from doing it any more - namely disconnect the user.

However, the ISP has a contract with the end user to supply a service. Can the ISP unilaterally cancel a service on just the say so of someone else. What legal right does the ISP have to do that, and what legal right does the 'someone' have to demand that the ISP does that?

Sure, the 'someone' can present whatever evidence they have to the ISP. But what ability does the ISP have to make a judgment on that 'evidence'? AFAIK, in Australia it requires a court to make judgments on those sort of things (in fact, right at the moment we are spending a lot of money on legal advice to find out).

But why is the ISP in the middle anyway?

If someone is drunk driving, no one goes to the MRD and demands they stop that driver from using the roads any more - with the implication that if the MRD doesn't take action, they somehow also share the blame.

If someone, say, plots a bank robbery with their nefarious mates using mobile phones, no one goes to the phone company and implies they are an accessory to the crime for not preventing the network from being used that way.

Yet those are the outrageous implications made in infringement notices sent to ISP's.

And, it seems, until it is tested in a court of law (at a cost of most likely not less than \$500,000 for the ISP), no one will know what the actual case will be for ISP's.

However, look at it from the agents acting for the copyright owners point of view. They see someone that they believe is stealing their clients property, and it is their job to act in their clients interest and stop that. How can they? All they have is an IP address that, when they look up who it belongs to, tells them it is some ISP somewhere in the world.

What are they to do? Write to the ISP and request the contact details of the user of the IP address? Ok, but in any country with privacy laws, they will most likely either be ignored, or at best, told that the information they seek can only be obtained with a warrant (which I believe is the case in Australia).

Then where does that leave them? Pretty much they have to say that 'well, we have told you we think one of your customers is stealing our clients property, you wont tell us who that person is or do anything about it, so we have to assume you (the ISP) is complicit in that theft.

Nice one. The ISP can't do anything because is breaches their contract with the customer and privacy laws. The Copyright agent can't protect their client because the ISP wont (or legally can't) give them any information or take action to further prevent it.

One thing is for sure; the lawyers on both sides can start booking their next family vacation in the Swiss Alps.

But maybe there is a solution. A pretty simple one too.

The complaint is between the copyright agent and the end user. The ISP is only the 'meat in the middle' because privacy laws seem to say end user information must not be divulged.

But surely if you think someone is stealing from you, you have a right to know the identity of that person? In most cases, if all you had was a phone number, it would be a trivial exercise to look up the number and find the persons name and address in the phone book. Why should not the same be true for IP addresses? Why should an IP address be any more 'private' than, say, your street address? It's the same thing - globally unique, one is where you live in the world, the other is where you 'live' in the Internet.

I remember an old credo of the Internet 'Information wants to be free'. If that is indeed the case, then it can only work if it goes both ways.

So one solution is; let the copyright agents have the contact details of the person they have the complaint about. Such agents are always lawyers themselves - and hence officers of the court. Is there a difference in handing over end user contact information to them, rather than to, say, a debt collection agency?

Of course there is a difference, spelled out specifically in the Communications Act, as our own lawyers have just pointed out to us in response to our request to them for advice:

Section 276 of the Telecommunications Act provides:

276. Current eligible persons

(1) An eligible person must not disclose or use any information or document that:

(a) relates to:

- (i) the contents or substance of a communication that has been carried by a carrier or carriage service provider; or
- (ii) the contents or substance of a communication that is being carried by a carrier or carriage service provider (including a communication that has been collected or received by such a carrier or provider for carriage by it but has not been delivered by it); or
- (iii) carriage services supplied, or intended to be supplied, to another person by a carrier or carriage service provider; or
- (iv) the affairs or personal particulars (including any unlisted telephone number or any address) of another person; and

(b) comes to the person's knowledge, or into the person's possession:

- (i) if the person is a carrier or carriage service provider--in connection with the person's business as such a carrier or provider; or
- (ii) if the person is an employee of a carrier or carriage service provider--because the person is employed by the carrier or provider in connection with its business as such a carrier or provider; or
- (iii) if the person is a telecommunications contractor--in connection with the person's business as such a contractor; or
- (iv) if the person is an employee of a telecommunications contractor--because the person is employed by the contractor in connection with its business as such a contractor.

...

(3) A person who contravenes this section is guilty of an offence punishable on conviction by imprisonment for a term not exceeding 2 years.

Such conduct would also be likely to breach section 7 of the Telecommunications (Interception and Access) Act 1979.

7. Telecommunications not to be intercepted

(1) A person shall not:

- (a) intercept;
- (b) authorize, suffer or permit another person to intercept; or
- (c) do any act or thing that will enable him or her or another person to intercept; a communication passing over a telecommunications system.

Personally, as much as I don't like the idea of someone using the Exetel service to steal someone else's property, I much more don't want to go to jail for 2 years for giving that alleged thief's information to someone else.

So it seems what is needed is a streamlined and automated 'warrant' system, whereby the copyright agents 'swear and avow' (under some threat of terrible reprisal) as to the strict reason they want the information and the purpose they will put it to. Secure certifications means that such information could be sought and obtained from the ISP with a high degree of automation - if not in real time, then pretty close to it. And the very great benefit, to both parties, is that from that point on they are in direct communication with each other, without, necessary legal, interruption of a third party ISP.

Posted by Steve Waddington at 10:50

Wednesday, July 30, 2008

Nine Interesting Things

John asked me 'What value did you get from the conference last week?'

As I mentioned, the event exceeded my expectations. I was expecting to maybe see one or two things 'of interest', and maybe pick up another one or two points of some value from the presentations. But there was far more than that.

After a review of my notes, the nine things of greatest value were:

- 1) Allot recognize the issue of mis-classification and have two ways it can be fixed. One is by turning off the classification memory, the second is to upgrade the service controller server.
- 2) It was clear that every service provider, including major regional Telco's and every major Australian service provider are using DPI of one sort or another to manipulate traffic
- 3) From Taiwan Telecom I found out an different way P2P policies can be set up in a network using TOS settings
- 4) There will be a 2-3 year period before DPI, at least from Allot, becomes a 'mainstream' product - however, they are probably further ahead or equal with other DPI companies
- 5) DPI alone is not an optimum traffic control solution, DPI + Caching leverages both technologies for much better results
- 6) Allot have released a new chassis based server, that can include a cache blade and also an intelligent network protection blade
- 7) Allot have acquired a company that makes a network protection server - it detects things like spam and bot-nets from inside the network and can automatically isolate and inform the end user. It also detects attacks from outside the network and can shut them down
- 8) We will almost certainly need to move to 10Gbps in our network core within the next 12 months
- 9) There is a low end Allot product we can use for Brisbane and Melbourne, and there can be unified policy control via the service controller server

Plus it was very interesting listening to what companies like NTT, China Telecom and Korea Telecom had to say, as well as the other 'smaller' (but huge by Australian standards, having more subscribers than our entire population) service providers from India, Indonesia and elsewhere. And interesting to see that as an 'early adopter' we are ahead of almost everyone, if not everyone, in the way we have deployed and integrated DPI and caching into our network.

Posted by Steve Waddington at 12:21

Tuesday, July 29, 2008

Scammed in Bangkok

Between the conference ending and our flight home, we had a free afternoon in Bangkok to do some sightseeing (read; 'shopping'). Before consigning myself to the roll of bag carrier for the afternoon, I particularly wanted to visit a gem shop that had been recommended to me by friends who have brought some very lovely sapphires there.

Looking at the tourist map of the city, the address we were heading for was about two blocks north of the Shangri-La Hotel - an easy landmark to find. Never the less, we first stopped at the concierge desk to just check (read; allay the fears of my travel companion) that my infallible sense of navigation in foreign cities was in fact infallible.

'Yes' the nice lady at the desk told us, after looking at the piece of paper with the address 'I know that street, just take the hotel shuttle boat to the MRT landing and it is about 15 minutes walk from there, or you can get a taxi'.

We got to the landing and I was looking around for a way to cross under the sky rail to get to the Shangri-La, from where I would better know the way - not the quickest route to our destination, but one I could be sure of.

While doing so, no doubt looking like clueless tourists, a friendly lady came up to us and said something like 'Hello, can I help you, I work at the hotel (the Hilton where we were staying) and I recognize you from there.'

Barbara started talking to her and they chatted for a minute about where we were from...

'Australia - Sydney?'

'No, Perth'

'Ah, my sister lives in Perth, she married to an Australian there, I go there next year'

... while I was trying to refold the map to some manageable size.

'Where you looking to go (sorry my English not so good)?' She asked

'We are looking to go to this address' I said, showing her the piece of paper.

'Ah, you looking to buy gems?' She asked, obviously recognizing the street or area where all the gem shops are.

What followed was about 10 minutes of very polite conversation of where to get the best gems, her showing us a beautiful star sapphire ring she was wearing and some other jewelry she had on. She went on to explain that where we were headed was really only for tourists, and the prices were higher. But she brought her gems from a government wholesale outlet, much cheaper. What is more, the quality is strictly controlled and guaranteed by the government.

'Well, we are going to have lunch first, and then we may go there if we get time' I said, it being about 1:30pm

She went on to comment that today, for the King's birthday, the government was excluding sales tax on gems. But that the government wholesaler would only be open until 3pm, and a Tuk Tuk could take us straight there, only 20 minutes away, but too far to walk, only cost 50 Baht. She then proceeded to write down the address for us, somewhere along the way asking if we spoke Thai. No, we didn't. so she kindly also wrote the directions in Thai that we could show the Tuk Tuk driver.

We thanked her and went on our way - toward the Shangri-La to recommence my navigation to the place we originally wanted to go.

Not more than 20 paces on, as we were heading towards a footbridge that looked like it would lead to the street the hotel was on, I must have showed some further clueless tourist hesitation, because a well dressed gentleman came up to us and asked where we were heading.

I told him we were looking to get some lunch. He hesitated a moment, and by way of friendly introduction told us he

was a doctor and his office was just nearby. 'Nowhere good here to eat' he told us 'You need to go into city for best food, come with me, I show you way'.

The way turned out to be a taxi rank nearby, where our newly acquired guide told us he would tell a Tuk Tuk driver to take us to where all the nicest food was. He was a dentist (a doctor dentist apparently), so he knew about good food. While he was instructing the Tuk Tuk driver where to take us, I told him we were meeting friends anyway, but thank you, and very nice to have met you, and we walked off.

Barbara asked me why we didn't take the tuk tuk to get some food (it was now about 2pm and we had both skipped breakfast and so were pretty hungry). 'My dear' I explained 'We were 100% being scammed'.

She didn't believe it. And why would she. Thai people are very friendly and commonly will help out a foreigner who looks a bit lost. The first lady who spoke to us could well have been from the hotel. She was nothing but friendly, polite and at no time was there any pressure, other than her hospitality, to follow her suggestion. So too was the doctor/dentist very polite, well dressed and well spoken. There was no apparent connection between him and the lady, they were just two people trying to be friendly and help out some tourists. Everything they said was completely plausible.

Never the less, with my now slightly grumpier companion (the address we were looking for was both farther and more difficult to get to, in the sweltering midday tropical heat, and we never did end up having lunch), we eventually made it to the address I wanted. And, incidentally got a fantastic price for about nine carats of gorgeous sapphires.

Still discussing on the way back to the hotel why I had declined the advice of those very nice people, I cited as reasons the coincidence of the lady just happening to know the 'best place for gems', the coincidence that, after asking where we were from, she just happened to have a sister who lived there. The coincidence that there was a special discount for the Kings birthday, but we had only until 3pm (not enough time to go back to the hotel and look it up on the Internet). The 'safety' that the gems were government guaranteed (when in fact there is no such thing). The way the lady made sure she asked us if we spoke Thai before she wrote, what could have been the directions in Thai. The 'doctors' slight hesitation when we said we were looking for a place to eat, rather than how to get to the government gem outlet. And then how he became a dentist, and steered us to the taxi rank with the one, conveniently waiting tuk tuk.

But all that was in vain. In the end I had to resort to showing her the web sites that explained exactly how such scams work - almost word for word, that I had read before we left Australia. Just google 'sapphire scam', there are at least thirty sites that expose what can happen.

Quite frankly, to me, the whole event was highly entertaining - in fact, the highlight of the trip. The sting was almost flawless, well executed, and completely plausible. You just can't pay for entertainment like that.

Posted by Steve Waddington at 12:32

Monday, July 28, 2008

The Summit Meeting

Henry Kissinger wasn't there, but just about everyone else was. I am talking about the 'Allot Broadband Service Optimization Summit 2008'. Allot of course make the NetEnforcer product that does deep packet inspection, analysis and control of IP traffic.

I have to say that the event was exceptionally well run, far exceeding my expectations in every respect. Not the least of which was the amount of information packed into the, essentially, two day event. First of all, the list of attendees was quite impressive, with representatives from telco's like China Telecom, NTT, Thai Telecom, TNZ, Singtel - in fact all the largest telco's in the Asia Pacific region were there.

Another interesting thing was the Australian ISP's who were, and were not there. Including one that was 'incognito', apparently, I can only guess they don't want anyone to know they are manipulating traffic. Talk _is_ cheap, and 'high performance internet' extends only as far as your DPI box, right guys?

In casual conversation over breakfast I was chatting with well connected industry person about the notable absence from the Aus ISP community. Not because they don't have DPI and traffic shaping, I was told, but because they have bought a competitors (to Allot) product (and I place some veracity on that information, because the person I was talking to has been marketing Allot to those companies and is in a position to know what they have).

So that is a long way to get to the point, which is - EVERY Telco and ISP (of any note) is, right now, doing DPI and traffic shaping. Whether they tell their customers or not.

The second great thing about the summit was the presentations. Very senior people from seriously serious Telcos, spoke about their DPI deployment, issues and experiences. Needless to say, a common concern expressed by everyone was the 'P2P Problem', which is twofold.

One; it can saturate even the largest most robust network.

Two; it is by far the cheapest way to distribute content. How cheap? I'm glad you asked. These figures are from Equinix to compare the cost of methods of content distribution (by the content owner) for 10 feature length videos (I assume the cost is US centric):

1. Transit (just buying enough bandwidth).....\$0.60c
2. Content Distributed Network (Like Akamai).....\$0.77c
3. Transit and CDN hybrid.....\$0.69c
4. P2P.....\$0.18c

Cheaper by a factor of at least three. As you would expect, because the end user is the one who 'pays' for the distribution bandwidth. All the distributor needs is something popular enough to inspire seeding. I can only imagine someone like Microsoft would save an absolute bundle on their massive server farms and bandwidth if, say, XP updates were P2P seeded. But that is something Linux distro people have known for a long time.

Being the cheapest distribution method isn't a problem to the distributor of course. But what it means to the ISP is that P2P traffic is only going to increase, and increase dramatically according to all the nice graphs analysts show, in the next five years. Further taxing networks and bandwidth. And while bandwidth can always be increased, the expense and time of 'raw' bandwidth just won't keep up with demand.

The solution (or perhaps selling point) is, surprisingly, not just DPI and shaping. Because what that does is improve the QOS, but at the expense of degrading the 'Quality of Experience' - QOE. What is needed therefore is a way to manage network bandwidth and protect it from P2P saturation, while at the same time delivering good 'QOE' for access to ever increasing high bandwidth content.

Enter the combination of DPI and P2P Caching.

No surprise that Allot and PeerApp recently announced a joint solution for Video optimization.

The great thing about that for Exetel, is that we will be able to leverage what we started as a 'go it alone' solution nine months ago with refinements Allot and PeerApp develop for the global market.

Posted by Steve Waddington at 11:20

Saturday, July 26. 2008

The Shape of Things to Come

In fact, the shape of things already here... if you live in Korea.

What KT are doing now with IPTV makes the latest Foxtel Digital offering look very outdated.

So what makes it so impressive? Here we go:

- Currently 68% of KT broadband customers are connected at 50Mbps via VSDL or FTTH. It will be 92% by 2010 (about the time the tender to begin to plan to think about starting to plan for the far more basic FTTN comes out in Australia)
- Providing capacity for UNLIMITED channels via multicast IP
- view any program any time
- Personalised services; essentially, every subscriber can create their own customised TV channel
- fully interactive services on the screen, including; news (a bit like google news - with the topics of interest to the subscriber), SMS style messaging in message pop up on screen, phone callse and CID pop up on the screen, share trading, shopping (with fully integrated product and best deal search)
- search engine like search for programs, news items, products, information etc.
- custom services for businesses such as in company broadcasting
- home media IP hub; service not just available on one or two TV screens, but to any networked device in the house
- Pop up information about the actor, or the car he/she drives, or the product he/she is using in the program
- real time on screen polling and feedback

(of course, just imagine the information that gives marketers)

- tracking of viewing habits and preferences to make intelligent suggestions of programs you might like to watch

And the absolute coup de gras killer app.....

- Watch sport from the view you select (changeable any time via a sidebar on the screen), or from the view point of your favourite player, or follow your favorite player through the match.

Whew indeed, if the 'whistlers' in the Foxtel ad saw what KT are doing, their lips would fall off.

Posted by Steve Waddington at 10:32

Tuesday, July 22. 2008

How stupid would you have to be...

To smoke? When the government goes to such great lengths to tell us how bad it is for us. But hey, who can really trust what the government says.

To drink and drive? If I just take it easy through the back streets, I'll be right. Besides, I'm too wasted to walk.

To blow your rent and food money at the Casino? Well, I could get lucky.

To not notice your boogie board had suddenly doubled in weight and volume. I er, well, er, umm. Baggage handlers done it!

To steal copyright material and/or distribute it over the Internet. Ah, now, there I am safe, because on the Internet I am completely anonymous and no one knows what I am doing, so I can never get caught.

Of those five statements, it's hard to pick which one is stupidest.

Yet, unbelievably, people will argue that sites with names like 'piratebay' and 'istole-it' are only used for completely legitimate reasons. Of course they are, and you are as pure as the driven snow, right?

Sure, those sites themselves contain none of the, irrefutably, copyright violating material their web sites list as available. We all know how p2p works. But it is ridiculously easy to track and log and identify the IP addresses of anyone who accesses, shares or redistributes that material - through exactly the same trackers that make it available in the first place.

Say you are the sort of person who, seeing a \$10 note fall unnoticed from someones pocket, have a look around, see you are unobserved, stoop and quickly pocket it yourself. Ha! No one saw. The perfect crime.

Say you are that same person, but, your quick glance around reveals; three people are looking at you, and two of them are policemen, you, for some reason, are wearing a sign with your name and address on it; and, oh yes, there is that big spotlight right on you.

"Excuse me Sir, you seem to have dropped a \$10 note" you say with a smile, having done your civic duty.

And yet it seems there are some stupid, stupid people, who must think that somehow, no one can see what they are doing just because they are on the Internet and not walking in the park.

How many people? Looking at the copyright infringement notices stats for the last week:

```
date notif block
20080714 3 0
20080715 21 7
20080716 17 3
20080718 8 4
20080719 6 2
20080720 9 0
20080721 7 3
```

About 10 per day, out of ~55,000. Roughly 1/5th of 1 percent. If you take that a percentile on an IQ bell curve. what is that, three sigma's out? That's not just stupid, that's borderline between mild and moderate mental retardation.

Posted by Steve Waddington at 11:09

Monday, July 21, 2008

Allot Conference

I am attending the 'Allot Broadband Service Optimization Summit' over the next few days - a grand sounding name to be sure. What I am looking to get from the trip is three things:

1. If it is possible to, and how, we can better optimize our NetEnforcer equipment, particularly with the PeerApp cache engine
2. The plans Allot have to improve the traffic classification algorithms
3. A review of why we seem to have so many traffic classification issues - having been told we are the only ones reporting it, and how we can fix that.

I have requested one-on-one meetings with their lead developer and their Director of Technology, so I am looking forward to some sensible discussions with them.

Perhaps it is serendipitous, but four weeks ago we were approached by Cisco to evaluate their Service Controller engine. and they were kind enough to give us a full briefing of its capabilities. So a secondary agenda will also be to get a feel of how whatever goodies Allot are releasing might stack up to the SCE.

It looks like Cisco are fairly serious about getting into that market. Some of the phrases passed across the table were things like 'very keen to do business' and 'trade in existing 3rd party equipment'. I am sure Allot are aware of that, so it will be interesting to see what their plans are.

Posted by Steve Waddington at 16:29

Thursday, July 17, 2008

The Unluckiest Carrier in the World...

... would have to be Optus.

<http://www.australianit.news.com.au/story/0,24897,24033460-15306,00.html>

Both of the circuits that provide full redundancy for the network cut, I mean, what are the chances of that? It really is bad luck. I guess sometimes these things just happen.

Then there was the time a couple of years ago (can't remember exactly when, or the fine details now), when a similar thing happened to their IP network - both circuits that provide redundancy cut, at the same time. One by a landslide/mudslide and the other by something else I have forgotten now (a backhoe maybe?).

And prior to that, going way back to the late 90's now, I recall a similar occurrence with International IP (both routers in different locations US failed at the same time, as I recall). And another time when something caused loss of the data link between Melbourne and Sydney for, was it 36 hours?, I don't recall, long enough to cause much consternation, anyway.

Still, four instances spread over 10 years, who can really complain in the long run? But what appallingly bad luck that both of the redundant circuits/equipment would fail at the same time on each occasion.

Posted by Steve Waddington at 16:30

Monday, July 14, 2008

Tree Planting Day

The beneficiary of Exetel's carbon neutralization program is Men of The Trees (MOTT). They have lots of information about the work they do on their web site; in a nutshell, they replant areas of degraded or salt affected land with many varieties of trees native to the area. It is worthwhile reading Barrie Oldfield's articles on the MOTT site for a full understanding of the programs they run.

I liked the MOTT programs during my initial investigation as to how Exetel could offset our carbon footprint because of the holistic approach they take. That is, there is more to offsetting carbon emissions than just turning it into wood, which I hope the rest of this blog article will demonstrate.

To put that in perspective, about two years ago I flew down to Albany (my home town). Over the last 25 years farming land in that area has been progressively 'put under trees', as farmers have leased out their land for the purpose of growing trees, as a commercial crop, both for carbon credits and for eventual harvest for woodchips and pulp.

Having flown over the area as a child with my Uncle in light aircraft, I still had a clear memory of how the landscape looked from the air. So flying in on the Skywest plane it was something of a surprise to see just how much of what used to be green cattle and sheep pasture was now blue/green trees - Tasmanian blue gums I believe they are.

What really jarred however was the way the paler blue gums stood out against the patches of darker native bush. Having grown up, hiked through, camped in and driven all around the south west, those trees just looked out of place to me.

That was the first thing. The second is the impact of mono culture tree farming, and using trees that are best for the eventual harvesting purpose. As I understand it, there is essentially not much difference between a tree crop and, say, a wheat crop, just a much longer time frame between harvests.

To stop insects attacking the trees, pesticides are used. To stop 'weeds' (in this case the normal undergrowth you would expect in a forest) and reduce fire risk, herbicides are used. Like any high concentration of a single species, disease is a threat, so more sprays are used to prevent that.

All in all, the cultivation needed for commercial tree farming can only cut into the overall effectiveness of tree farming has in offsetting environmental impact. Nor can it do the native fauna much good - which are just as unwelcome on a tree farm as they hopping, pecking or jumping their way through a wheat crop.

So you can maybe see what is appealing about the Men of the Trees programs. They replant using either local native trees, or species close to it (in some cases where more salt tolerance is needed) and they use typically 15 to 20 different species in each area.

Apart from the initial spray to prepare the re-plant area (to kill off the thistle, dock and cape weed), as far as I know, no other chemicals are used from then on. The flow-on effect is not hard to imagine - native trees promote a healthy under story, making new 'havens' for native fauna, and the return of other native flora species. Once planted there is no more cultivation needed, which must make for more effective carbon offset.

Of course, it isn't a commercial crop - it is done purely for the environmental benefit. While there is no direct commercial return for the farmer, there is however almost certainly quite a few indirect benefits. Here is a brief synopsis of my own first hand experience:

- the farmer has an erosion or salt problem
- he asks MOTT to replant the affected area
- MOTT undertake the project and provide the labour and the plant stock
- the farmer does the initial preparation and, if needed, fences off the areas
- The MOTT condition for replanting is a 100 year caveat on the land
- soil degradation is stopped and reversed
- topsoil run-off is stopped
- the water table is lowered, reducing salinity

MOTT planting is done over the winter months, by MOTT staff and volunteers. This planting season, I volunteered my own and my families services to help out with the planting. For two reasons 1) as corporate sponsors (and the only Exetel staff in WA), it seemed like the appropriate thing to do, and 2) as corporate sponsors, reading about it on a web site is one thing, but seeing, and participating first hand, is what I really wanted to do.

And so, early on a rainy Wednesday morning Barbara and I roused the kids, loaded the thermos and sandwiches and headed up towards York for a MOTT arranged 'corporate planting day'.

Men of the Trees assembly area, but where are all the men?

73,000 trees as tube stock

Loading up, ready for planting.

Area prepared by the farmer for tree planting. This is about 1/8th of the total area planted on the day - 500 trees will go here.

The only thing more embarrassing than getting a tractor like this bogged, is getting your neighbors tractor bogged trying to pull it out

How to use the 'potting foot' device for planting trees - 1. Shove the sharp end into the ground (and not your foot)

How to use the 'potting foot' device for planting trees - 2. Put your foot on the lever to open the jaws and make a hole, then drop the tube stock into the, er, tube.

How to use the 'potting foot' device for planting trees - 3. Observe the tree in the perfect hole you have made

How to use the 'potting foot' device for planting trees - 4. Cave in the hole with your foot. 5. Repeat many, many times.

In five years time, what was once prime cape weed growing land will be a new little forest of native trees.

The combination of winter planting, using the potting foot tubes, very good nursery tube stock and appropriate plant selection for the area equates to an 80% or better survival rate of the planted trees.

Working the daylight hours we planted out 4,500 trees - a very satisfying days work.

Posted by Steve Waddington at 09:34

Friday, July 4, 2008

The end of File Sharing?

When I was very young, maybe five or six, I remember clearly driving home with my parents from my grandmothers one Sunday night, and them commenting on which other drivers were 'on their way home from the session' as we made our way along Albany Highway the few miles to our house.

The Sunday session was an institution in WA for many years; where the hotels and pubs could only open for a couple of hours on Sunday, and so the mission was that an afternoon and evenings drinking would be accommodated within those few hours. Booze buses and RBT was still a thing of the future in those days, and while still of course illegal, the penalties for drink driving were far milder, and the chances of being picked up even less. The fact was, it was common place to drink and drive. Everyone one knew it was wrong (amazing people could come to that conclusion without the millions spent since on 'bloody idiot' ad campaigns, but there you go), and just how wrong could it be when a) hardly anyone got caught, b) when they did, the penalty was only a fine of a few quid and c) most everyone did it anyway?

Because of course, if everyone else does it, it must be alright, right?

Then, over time, governments have changed legislation, the penalties have become far steeper and the police far more pro-active policing that particular anti-social crime. Anyone who contemplates ding driving now knows, for sure, they will get caught (maybe not today, but one day), and, quite rightly, will be facing criminal charges, steep fines, and loss of license.

I think the parallel with downloading copyright material is obvious. Obviously not physically, but certainly in terms of policing, penalties and general public acceptance, because:

- a) as far as I know, in Australia, no one has been 'caught' (action was taken against one ISP, but settled out of court)
- b) so who knows what the penalties are - but 'zero' at this stage
- c) apparently, at least judging by the plaintive cries from self confessed offenders, 'everyone does it'

To say 'everyone does it' is without doubt a gross overstatement made by people with a fairly small circle of like minded acquaintances. Never the less, it is safe to say that, while in the minority, there are people who otherwise see themselves as law abiding citizens but frequently, perhaps daily, share copyright material - that is; engage in theft - and do not think there is anything wrong with that.

What is needed is a massive public education campaign to let 'everyone' know that file sharing of copyright material is wrong. If you file share, you're a bloody idiot, you have a small penis and the top of your skull will be removed and everyone will be able to see your brain. That will put an end to it once and for all. Yes, of course it will Virginia.

Moving right along.

What will stop file sharing dead is, just like drink driving, the certain knowledge of detection followed just as certainly by prosecution. The reality of a \$20,000 fine and/or six months in prison would soon reduce 'everyone' to either the terminally stupid or the criminal hardcore (though it is hard to imagine a real criminal taking that sort of risk to save \$5 on video rental).

Detection is no problem. In the past notices have arrived from mostly US based organisations representing the artists rights. Perhaps the perception is that the US is a long way away, and their claim under international copyright law doesn't apply here. But then today, both via email and CD Hard copy the Australian Federation Against Copyright Theft has stepped in. This is the alleged infringement list they sent us:

Peer IP	Date and Time UTC	File Name Downloaded
xxx.xxx.10.196	27-06-2008 10:42	Transformers[2007]DvDrip[Eng]-aXXo
xxx.xxx.105.158	25-06-2008 19:25	Hitman[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.109.80	29-06-2008 01:43	Ugly.Betty.S02E17.HDTV.XViD-DOT.avi
xxx.xxx.109.80	29-06-2008 01:07	Ugly.Betty.S02E18.HDTV.XViD-DOT.avi
xxx.xxx.11.21	23-06-2008 11:56	Lost.S04E10.PROPER.HDTV.XViD-DOT.avi
xxx.xxx.11.21	26-06-2008 13:12	Lost.4x12.Theres_No_Place_Like_Home.HDTV_XviD-FoV.avi

xxx.xxx.110.229 28-06-2008 06:04 Lost.S04E13-E14.HDTV.XviD-2HD.avi
xxx.xxx.110.77 25-06-2008 12:07 Evan.Almighty[2007]DvDrip[Eng]-aXXo
xxx.xxx.110.77 25-06-2008 11:48 Spider-Man.3[2007]DvDrip[Eng]-aXXo
xxx.xxx.110.77 25-06-2008 15:30 American.Gangster[2007][Unrated.Edition]DvDrip[Eng]-FXG
xxx.xxx.112.130 29-06-2008 09:46 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.114.131 26-06-2008 15:38 Supernatural.S03E15.HDTV.XviD-0TV.avi
xxx.xxx.115.112 26-06-2008 06:52 Spider-Man.3[2007]DvDrip[Eng]-aXXo
xxx.xxx.116.183 29-06-2008 11:00 Fantastic.Four-Rise.Of.The.Silver.Surfer[2007]DvDrip.AC3[Eng]-aXXo
xxx.xxx.119.56 27-06-2008 17:46 Transformers[2007]DvDrip[Eng]-aXXo
xxx.xxx.119.56 28-06-2008 10:37 28.Weeks.Later[2007]DvDrip[Eng]-aXXo
xxx.xxx.120.115 24-06-2008 16:23 Shooter[2007]DvDrip[Eng]-aXXo
xxx.xxx.120.237 25-06-2008 18:20 Hitman[2007][Unrated.Edition]DvDrip[Eng]-FXG
xxx.xxx.122.27 23-06-2008 09:23 Greys.Anatomy.S04E15.HDTV.XViD-DOT.avi
xxx.xxx.122.27 23-06-2008 10:42 Greys.Anatomy.S04E14.HDTV.XviD-2HD.avi
xxx.xxx.122.40 29-06-2008 10:00 Ratatouille[2007]DvDrip[Eng]-aXXo
xxx.xxx.122.40 24-06-2008 20:04 National.Treasure.2-Book.Of.Secrets[2007]DvDrip[Eng]-aXXo
xxx.xxx.123.65 24-06-2008 06:22 Evan.Almighty[2007]DvDrip[Eng]-aXXo
xxx.xxx.124.236 29-06-2008 16:38 Lost.4x12.Theres_No_Place_Like_Home.HDTV_XviD-FoV.avi
xxx.xxx.124.236 29-06-2008 12:55 Lost.S04E10.PROPER.HDTV.XViD-DOT.avi
xxx.xxx.160.118 26-06-2008 14:03 Bones.S03E13.HDTV.XviD-NoTV.avi
xxx.xxx.160.118 24-06-2008 14:30 Bones.S03E14.HDTV.XviD-0TV.avi
xxx.xxx.163.250 28-06-2008 14:39 The.Chronicles.of.Narnia.Prince.Caspian.2008.Eng.TS.DivX-LTT
xxx.xxx.163.49 25-06-2008 05:10 Supernatural.S03E15.HDTV.XviD-0TV.avi
xxx.xxx.164.243 24-06-2008 00:39 American.Dad.S03E16.PDTV.XviD-0TV.avi
xxx.xxx.166.73 23-06-2008 08:06 Lost.S04E11.HDTV.XviD-2HD.avi
xxx.xxx.169.119 27-06-2008 16:33 Samantha.Who.S01E15.HDTV.XviD-XOR.avi
xxx.xxx.169.143 24-06-2008 07:09 Two.and.a.Half.Men.S05E16.HDTV.XviD-XOR.avi
xxx.xxx.169.166 24-06-2008 00:25 Two.and.a.Half.Men.S05E16.HDTV.XviD-XOR.avi
xxx.xxx.17.88 24-06-2008 07:08 Lost.S04E11.HDTV.XviD-2HD.avi
xxx.xxx.17.88 29-06-2008 20:49 Lost.4x12.Theres_No_Place_Like_Home.HDTV_XviD-FoV.avi
xxx.xxx.173.126 27-06-2008 00:07 Casino.Royale[2006]DvDrip[Eng]-aXXo
xxx.xxx.173.34 24-06-2008 04:17 The.Big.Bang.Theory.S01E17.HDTV.XviD-LOL.avi
xxx.xxx.173.55 25-06-2008 15:03 Blood.Diamond[2006]DvDrip[Eng]-aXXo
xxx.xxx.176.102 29-06-2008 22:37 Stardust[2007]DvDrip[Eng]-FXG
xxx.xxx.178.146 25-06-2008 22:03 War.Of.The.Worlds.(2005).[ENG].[DVDrip].[WS]
xxx.xxx.178.23 29-06-2008 03:16 Lost.S04E13-E14.HDTV.XviD-2HD.avi
xxx.xxx.178.56 29-06-2008 12:22 Casino.Royale[2006]DvDrip[Eng]-aXXo
xxx.xxx.181.34 28-06-2008 11:26 Greys.Anatomy.S04E14.HDTV.XviD-2HD.avi
xxx.xxx.182.173 26-06-2008 23:58 American.Dad.S03E13.PDTV.XviD-XOR.avi
xxx.xxx.182.173 28-06-2008 10:01 American.Dad.S03E16.PDTV.XviD-0TV.avi
xxx.xxx.182.173 28-06-2008 11:24 The.Simpsons.S19E20.PROPER.PDTV.XviD-E7.avi
xxx.xxx.183.61 24-06-2008 10:08 Lost.S04E11.HDTV.XviD-2HD.avi
xxx.xxx.186.32 29-06-2008 11:23 Family.Guy.S06E12.PDTV.XviD-2HD.avi
xxx.xxx.187.119 29-06-2008 04:47 Lost.S04E13-E14.HDTV.XviD-2HD.avi
xxx.xxx.187.119 25-06-2008 03:38 Lost.4x12.Theres_No_Place_Like_Home.HDTV_XviD-FoV.avi
xxx.xxx.187.119 29-06-2008 20:38 The.Chronicles.of.Narnia.Prince.Caspian.2008.Eng.TS.DivX-LTT
xxx.xxx.188.200 28-06-2008 12:06 The.Big.Bang.Theory.S01E17.HDTV.XviD-LOL.avi
xxx.xxx.19.72 25-06-2008 07:27 The.Big.Bang.Theory.S01E16.HDTV.XviD-XOR.avi
xxx.xxx.190.170 24-06-2008 12:57 Hitman[2007][Unrated.Edition]DvDrip[Eng]-FXG
xxx.xxx.190.209 28-06-2008 14:21 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.191.125 26-06-2008 09:40 300[2006]DvDrip[Eng]-aXXo
xxx.xxx.191.90 29-06-2008 11:06 Supernatural.S03E13.HDTV.XviD-HDQ.avi
xxx.xxx.191.90 29-06-2008 09:38 Supernatural.S03E15.HDTV.XviD-0TV.avi
xxx.xxx.192.51 26-06-2008 14:07 The.Prestige[2006]DvDrip[Eng]-aXXo
xxx.xxx.194.48 23-06-2008 07:49 The.Simpsons.S19E20.PROPER.PDTV.XviD-E7.avi
xxx.xxx.197.115 25-06-2008 10:11 The.Simpsons.S19E20.PROPER.PDTV.XviD-E7.avi
xxx.xxx.197.127 27-06-2008 22:52 Stardust[2007]DvDrip[Eng]-FXG
xxx.xxx.197.193 26-06-2008 00:04 Bones.S03E15.HDTV.XviD-LOL.avi
xxx.xxx.197.193 25-06-2008 01:01 Bones.S03E13.HDTV.XviD-NoTV.avi
xxx.xxx.197.86 27-06-2008 21:17 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.198.90 28-06-2008 09:43 American.Dad.S03E16.PDTV.XviD-0TV.avi

xxx.xxx.20.148 24-06-2008 20:22 Hitman[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.20.156 28-06-2008 23:26 Desperate.Housewives.S04E15.HDTV.XviD-0TV.avi
xxx.xxx.204.13 29-06-2008 13:31 American Gangster[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.208.153 29-06-2008 08:38 Night.At.The.Museum[2006]DvDrip[Eng]-aXXo
xxx.xxx.212.23 29-06-2008 02:11 The.Simpsons.S19E18.PDTV.XviD-2HD.avi
xxx.xxx.212.23 27-06-2008 15:33 The.Simpsons.S19E19.PDTV.XviD-2HD.avi
xxx.xxx.212.253 24-06-2008 09:05 The.Prestige[2006]DvDrip[Eng]-aXXo
xxx.xxx.213.232 23-06-2008 22:12 National.Treasure.2-Book.Of.Secrets[2007]DvDrip[Eng]-aXXo
xxx.xxx.215.46 23-06-2008 08:43 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.216.59 25-06-2008 13:49 Hot.Rod[2007]DvDrip.AC3[Eng]-aXXo
xxx.xxx.217.190 25-06-2008 13:17 Lost.S04E13-E14.HDTV.XviD-2HD.avi
xxx.xxx.219.139 26-06-2008 00:45 Ugly.Betty.S02E18.HDTV.XviD-DOT.avi
xxx.xxx.222.64 29-06-2008 01:04 The.Simpsons.S19E18.PROPER.PDTV.XviD-E7.avi
xxx.xxx.222.64 29-06-2008 02:26 The.Simpsons.S19E20.PROPER.PDTV.XviD-E7.avi
xxx.xxx.223.119 26-06-2008 09:27 Juno.2007.English.DVDSCR.DivX-LTT
xxx.xxx.223.241 28-06-2008 06:03 The.Simpsons.S19E20.PROPER.PDTV.XviD-E7.avi
xxx.xxx.226.1 24-06-2008 13:49 Casino.Royale[2006]DvDrip[Eng]-aXXo
xxx.xxx.226.117 25-06-2008 10:41 Ratatouille[2007]DvDrip[Eng]-aXXo
xxx.xxx.228.173 25-06-2008 09:54 Live.Free.Or.Die.Hard[2007]DvDrip[Eng]-aXXo
xxx.xxx.229.12 25-06-2008 13:11 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.229.253 25-06-2008 00:52 Hitman[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.23.184 29-06-2008 14:18 Evan.Almighty[2007]DvDrip[Eng]-aXXo
xxx.xxx.239.207 24-06-2008 08:46 Harry.Potter.And.The.Order.Of.The.Phoenix[2007]DvDrip[Eng]-aXXo
xxx.xxx.241.48 27-06-2008 18:49 Shark.S02E16.HDTV.XviD-LOL.avi
xxx.xxx.25.173 26-06-2008 09:27 Blood.Diamond[2006]DvDrip[Eng]-aXXo
xxx.xxx.25.173 27-06-2008 13:56 Evan.Almighty[2007]DvDrip[Eng]-aXXo
xxx.xxx.25.52 23-06-2008 18:10 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.26.207 26-06-2008 04:26 Supernatural.S03E13.HDTV.XviD-HDQ.avi
xxx.xxx.26.220 24-06-2008 09:33 American Gangster[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.28.124 28-06-2008 17:51 Flags.Of.Our.Fathers[2006]DvDrip[Eng]-aXXo
xxx.xxx.28.43 25-06-2008 06:58 Ugly.Betty.S02E16.HDTV.XviD-DOT.avi
xxx.xxx.28.43 25-06-2008 07:07 Ugly.Betty.S02E17.HDTV.XviD-DOT.avi
xxx.xxx.28.43 25-06-2008 06:22 Ugly.Betty.S02E18.HDTV.XviD-DOT.avi
xxx.xxx.29.161 23-06-2008 15:40 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.29.233 28-06-2008 15:12 Stardust[2007]DvDrip[Eng]-FXG
xxx.xxx.30.57 25-06-2008 17:39 Ugly.Betty.S02E18.HDTV.XviD-DOT.avi
xxx.xxx.4.154 29-06-2008 10:06 National.Treasure.2-Book.Of.Secrets[2007]DvDrip[Eng]-aXXo
xxx.xxx.4.241 29-06-2008 09:25 The Godfather part III
xxx.xxx.52.203 26-06-2008 17:56 Bones.S03E15.HDTV.XviD-LOL.avi
xxx.xxx.52.203 27-06-2008 09:48 Bones.S03E14.HDTV.XviD-0TV.avi
xxx.xxx.70.182 27-06-2008 10:14 Greys.Anatomy.S04E14.HDTV.XviD-2HD.avi
xxx.xxx.70.49 28-06-2008 09:40 Enchanted[2007]DvDrip[Eng]-FXG
xxx.xxx.70.49 28-06-2008 10:09 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.70.60 23-06-2008 05:58 Supernatural.S03E15.HDTV.XviD-0TV.avi
xxx.xxx.70.60 25-06-2008 09:13 Supernatural.S03E13.HDTV.XviD-HDQ.avi
xxx.xxx.73.119 26-06-2008 09:52 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.73.247 26-06-2008 01:27 National.Treasure.2-Book.Of.Secrets[2007]DvDrip[Eng]-aXXo
xxx.xxx.77.184 28-06-2008 15:50 Ugly.Betty.S02E16.HDTV.XviD-DOT.avi
xxx.xxx.79.79 26-06-2008 08:40 Hitman[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.8.198 29-06-2008 11:24 Atonement.2007.DvDRip.Eng-FxM
xxx.xxx.81.173 27-06-2008 21:42 Lost.4x12.Theres_No_Place_Like_Home.HDTV_XviD-FoV.avi
xxx.xxx.81.173 25-06-2008 21:43 American Gangster[2007][Unrated Edition]DvDrip[Eng]-FXG
xxx.xxx.81.71 25-06-2008 11:23 The.Prestige[2006]DvDrip[Eng]-aXXo
xxx.xxx.83.217 26-06-2008 08:26 Stardust[2007]DvDrip[Eng]-FXG
xxx.xxx.84.16 24-06-2008 14:30 Family.Guy.S06E11.PDTV.XviD-XOR.avi
xxx.xxx.84.253 28-06-2008 19:26 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.86.137 25-06-2008 16:57 Lost.S04E13-E14.HDTV.XviD-2HD.avi
xxx.xxx.88.108 28-06-2008 11:15 Evan.Almighty[2007]DvDrip[Eng]-aXXo
xxx.xxx.88.76 23-06-2008 07:50 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.88.90 23-06-2008 00:58 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.89.141 29-06-2008 18:30 Fantastic.Four-Rise.Of.The.Silver.Surfer[2007]DvDrip.AC3[Eng]-aXXo

xxx.xxx.9.83 28-06-2008 08:49 The.Kingdom[2007]DvDrip[Eng]-FXG
xxx.xxx.90.172 24-06-2008 12:38 Disturbia[2007]DvDrip[Eng]-aXXo
xxx.xxx.91.231 27-06-2008 11:10 Family.Guy.S06E11.PDTV.XviD-XOR.avi
xxx.xxx.92.94 28-06-2008 21:37 The.Spiderwick.Chronicles[2008]DvDrip-aXXo
xxx.xxx.93.45 24-06-2008 07:23 28.Weeks.Later[2007]DvDrip[Eng]-aXXo
xxx.xxx.94.187 23-06-2008 06:50 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.95.40 24-06-2008 10:53 American.Dad.S03E16.PDTV.XviD-0TV.avi
xxx.xxx.95.40 24-06-2008 10:28 Family.Guy.S06E11.PDTV.XviD-XOR.avi
xxx.xxx.35.160 28-06-2008 14:39 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.36.88 25-06-2008 23:35 A.Mighty.Heart[2007]DvDrip[Eng]-aXXo
xxx.xxx.39.53 25-06-2008 07:16 The.Big.Bang.Theory.S01E17.HDTV.XviD-LOL.avi
xxx.xxx.41.69 27-06-2008 12:46 Pirates.of.the.Caribbean-Dead.Man's.Chest[2006]DvDrip[Eng]-aXXo
xxx.xxx.42.81 23-06-2008 12:30 Greys_Anatomy.4x16_4x17.Freedom.HDTV_XviD-FoV.avi
xxx.xxx.42.81 23-06-2008 11:29 Greys.Anatomy.S04E15.HDTV.XviD-DOT.avi
xxx.xxx.70.190 23-06-2008 12:02 Greys.Anatomy.S04E14.HDTV.XviD-2HD.avi
xxx.xxx.72.110 29-06-2008 14:54 Juno.2007.English.DVDSCR.DivX-LTT
xxx.xxx.74.171 29-06-2008 19:26 Atonement.2007.DvDRip.Eng-FxM
xxx.xxx.75.110 23-06-2008 11:05 Supernatural.S03E15.HDTV.XviD-0TV.avi
xxx.xxx.75.110 28-06-2008 13:41 Supernatural.S03E13.HDTV.XviD-HDQ.avi
xxx.xxx.78.222 27-06-2008 19:54 The.Spiderwick.Chronicles[2008]DvDrip-aXXo

My guess would be that the only thing that can follow next is prosecution. Not today, maybe not for a while, but I would think the day can not be too far off.

If I were a file sharer, I would think now would be a real good time to start breaking myself from that nasty habit.

(aside: the notices sent to us by AFACT and other similar organizations indicate that 'everyone' is something less than 0.25% of the Exetel user base. At less than one quarter of one percent, maybe it is only the terminally stupid that file share copyright material after all.)

Posted by Steve Waddington at 06:38

Wednesday, July 2, 2008

SMS - The Best thing for business since...

... something else really good. Saying 'email' would be going too far, since, conceivably, one could live without SMS's, but what business could survive today without email? Maybe 'the stapler' would be a good analogy - not really missed if it's gone, but useful when it's there, and every home and business has one. But it's more than that, in fact, I would go as far as to say that SMS's will become just as integral to an organisation in the next few years as the fax machine is today.

For a long time now I have been getting SMS messages from the National Poker League telling me about poker games in Sydney. I can't remember when I gave them my mobile number, it must have been a few years ago. Since, occasionally, I like to play a hand or two of poker, I have never bothered sending 'STOP', and the once a fortnight frequency of the messages is no bother. So too in another life, John's idea at the time I think, did we send SMS messages to customers who wanted them if there was an ADSL outage. But letting customers/clients know about a particular event is just one use SMS lends itself to.

Just about anywhere you look where notifications need to be given, or timely information passed on, SMS's seem to fit the bill either:

- easier
- faster
- cheaper
- or more timely

than any other method.

Here are a couple of examples; My dentist calls me two to three weeks prior to an appointment to book a time for a check up. Fine, that part needs person to person contact. A day or two before, they call to confirm the appointment. The receptionist has to call, talk, confirm, leave a message, call back - whatever the many iterations are. And then I forget anyway. How much better if instead of a call, they send an SMS a couple of days prior, (for what a few cents?), and then another one on the day two hours prior. First, I have a text record of the appointment, so I can't misremember the date/time. Second, I get reminded on the day - all for what must be at least 1/10th of the cost and time a person making a phone call would be.

Today I received an SMS from a local Thai restaurant, telling me about a menu special. I must have dropped my business card in their 'win a free meal' jar sometime I guess. Every week my mailbox in North Sydney is stuffed full of flyer's from restaurants and cafes. Whatever specials they offer have to be at least weekly, and most of the time the volume in my mail box means the whole bundle goes straight in the bin. The cost of flyer's dropped to a mail box is what, \$50 per 1,000? Plus the cost of design, printing and delivery. And maybe one in twenty get read (I think that would be optimistic). An SMS has to be way cheaper, can be immediate (special 1/2 price menu tonight only!), and has a 100% chance of being read.

Then there are all the other personal uses (apart from between friends, family and lovers). This year my wife is managing my youngest son's soccer team, and uses sms's extensively to tell other parents about training (cancelled due to weather etc), game time changes etc. Which is so very much faster and more convenient than the 'telephone relay' method.

Our school uses SMS to advise if a child is away sick, or to update on an event pertaining to a particular year or class. Way more efficient than having to call, or be called by, the office admin staff.

So maybe I was wrong, and it is as good as email. Maybe already we can't live without SMS.

Posted by Steve Waddington at 13:28

Friday, June 20. 2008

Canberra

The Mk 20 saw service from 1956 to 1984 I think. And elegant aircraft, not unlike a pelican - funny looking on the ground, but very graceful in the air, I always thought.

It's an old lame joke from my Air Cadet days. Whenever 'Canberra' was mentioned in any context someone was bound to say something like 'the Mk20 or the Mk58?'.
Canberra

Anyway, Our co-lo rack space for the Canberra POP will be ready by the 3rd of July, the backhaul circuit on the same day, and the carrier cross-connect the day after on the 4th of July.

We haven't heard of any hold-ups from the carrier yet, and it is very unlikely that there wont be some, always 'surprising', issue that causes one. But can only cross that bridge when we come to it.

The 8th of July is the target commissioning date.

Posted by Steve Waddington at 11:30

Monday, June 16, 2008

Adelaide On Line

After all the cafuffle of getting the circuits in, the equipment installation went very smoothly.

I arrived at the Adelaide co-lo a few minutes before 8am on Friday, was greeted by the facilities manager there. Induction took 15 minutes or so. Equipment had arrived the day before, shipped from Sydney, along with all the patch cables, mounting kit and a couple of screwdrivers.

By 8:15 am it was all bolted into the rack. But uh oh, no patch cables had been shipped to connect it to the supplier circuits. Oh well, these things happen, I had allowed for that in my planning and fully expected to have to make a trip to an electrical and data supplier store to pick up whatever might have been missed.

I asked the facilities manager where one might be near by. His response 'I probably have those here, what do you need?'

"A single mode fibre patch lead and a cat 6 ethernet patch lead" I told him.

"Sure, here you go" he said, fishing them out of a box of new cable bits.

That was just so helpful, saving the half day running around I fully expected.

A few minutes later the cables were plugged in and the router was booting up.

The next inevitable 'hassle' is the loopback/activation of the carrier circuits. What happens is the carrier installs the circuit, which, before any equipment goes in is 'open' - there being nothing to terminate the circuit on at that stage. An open circuit causes an alarm on the carrier network, and so, quite rightly, the field tech puts a loopback on the circuit, to stop the alarm and to be removed when the customer equipment is ready.

So then the customer, us in this case, connects our equipment, and while the circuit is 'up', we are really only seeing the loopback. So too the carrier sees the loopback, so to them the circuit is also up, but of course nothing can work until the loopback is removed. What usually ensues is we make a call to the carrier to ask them to activate the circuit. The carrier says something like 'the circuit is activated, we can see it up'. Then there is a round of explaining that yes, it looks up, but that is the loopback, usually followed by the carrier saying 'so there is a fault then?', 'no, not a fault, we just want the loopback removed'. and eventually, often the next day, the loopback is removed and the circuit can work.

Convincing and cajoling to get the loopback(s) removed was the next part of my time allocation. My goal was to get it done before COB Friday, but if not, to escalate it far enough so that the work we needed to do once the circuits were up could be completed Saturday morning. I figured with the physical installation done well ahead of plan, I was looking good to get it all done on Friday.

By this time I had the laptop connected to the router and could ping each interface to see just what had IP connectivity and what didn't. The WAN circuit to Sydney pinged fine with just 22ms latency - which was good. Now the carrier circuit - hey, that can be pinged too!

Oh, wait a minute, I am probably pinging our router interface IP, not the carrier interface. Hmm, that IP pings too, and I was right the first time. It can't possibly be that easy, what have I missed?

I checked and checked again, but nothing was missing. It was now 8:50am and the whole lot was working. Just amazing.

Vinna completed the final software config on the router, and by noon we were able to announce the Adelaide POP 'in production'.

The only downside was I couldn't find any of the free wifi hotspots in Adelaide. There were plenty of hotspots, but all credit card or sms subscription. Expensive if you need to spend a working day (as I now did, having nothing else productive to do) on email. So I ended up paying the gouging hotel rate for 24 hour 'broadband'.

Posted by Steve Waddington at 15:30

Thursday, June 5, 2008

Perth Live

The first ADSL circuit, my home connection, was cut over to our Perth LNS on Tuesday. At the moment I have the luxury of being the only occupant of the LNS router and all the bandwidth.

The change from Zedres to TW took place at about 10am, which I noticed when the circuit went dead. Logging on to my Cisco CPE I ran debug for PPP authentication and PPP negotiation, which showed that a) the DSL line was up and exchanging information with the Wellington Street TW LAC b) there was no authentication with our radius server past that point.

I asked Vinna to check that the Radius server could be reached by the Perth LNS, and it was. So we deduced the problem was somewhere within the carrier - which it was. At 12:55pm they flicked whatever switch they had previously forgotten to flick and the circuit completed negotiation almost immediately.

With that done, any future transfers/connections should take place with not more than a few minutes downtime for the end user.

The latency from the Perth POP to Sydney averages 63ms - a 10ms improvement over the previous circuit, no doubt because of the removal of the whole layer of carrier routers and backhaul equipment.

And so far there is 100% uptime for the POP - something that is unlikely to be improved on in the longer term.

Now that the main work for Perth is completed, we will be activating WAIX peering, WAIX, if you don't know, is the most active ISP peering exchange in Australia. It is a non-profit peering point, created for, and run by, the peer members (of which, Exetel will soon be one).

It is so popular that it is just about impossible to be a WA ISP without being a member of WAIX. I don't know what proportion of bandwidth WAIX will provide, but I would guess it will be about 5%. The important reason for using WAIX however is for gaming and other local content, but particularly gaming, where latency between other WAIX connected ISP's will be 10ms or less.

The likely completion date for the WAIX connection is this Friday.

Posted by Steve Waddington at 09:26

Wednesday, June 4, 2008

Adelaide POP II

The news for the Adelaide POP is disappointing. After my optimism last week that it could be commissioned this week, we have been told by the carrier that they now have to do a complete redesign.

The new date they gave us was July 4. However, pointing out a few facts to them - such as; the co-lo work was completed by the 30th, as expected; there is now less work for them to do as a result; nothing else in the design has changed; their delay is causing loss and disadvantage to Exetel in the market place. Has resulted in their new completion date being moved to June 6 or 'Monday at the latest'.

You know, for a moment there I thought this really was going to be the first carrier circuit delivered on time (ever). After 25 years, you think I would have known better. Hope springs eternal as they say.

In the meantime, we have now decommissioned the 2nd LNS at our Sydney POP, and that router will be used as the LNS for Adelaide. (An Adelaide POP will also encompass the Northern Territory as well as South Australia, so the only exceptions we will to national coverage will be Tasmania and Canberra).

Brendon has the router and associated bits ready to ship as soon as we have a firm date for circuit completion. As it looks at the moment, I may well be able to stop off and set it up en-route to Sydney next week.

Posted by Steve Waddington at 12:51

Friday, May 30, 2008

Adelaide POP

It seems to be a common theme that the last few metres are the hold up for the other 1,000 km of the circuit.

The ADSL wholesale supplier went to our Adelaide co-lo earlier in the week to install the cross-connect for ADSL, and were informed by the co-lo manager that they could only install it to the demarcation point, and that they (the co-lo provider) would have to run the cable from their to our rack.

That isn't really a problem - many co-lo's have a similar policy (although with Telstra it is often the exception that they let them run the cable directly to the customer rack). So we are now just waiting on the co-lo provider to complete their bit of the cabling and we should be ready to go. Assuming there are no more hold-ups, we could have the Adelaide POP running by as early as the end of next week. If that does happen, it will be the fastest setup for a POP I can remember - mainly because it is the first time, ever that I recall, that a carrier circuit has met the SFOA/eCRA 30 day delivery time.

What is always annoying about these type of installations though (and it never seems to change), is the way a standard connection comes as such a surprise to everyone involved. It is like the circuit and/or service ordered is the very first one any of the suppliers in the chain have every done, and it all comes to a complete full stop as soon as any irregularity, no matter how minor, is encountered.

Like the above just for example. It is common practice for circuits in co-lo's to be delivered to a demarcation point. Any installation technician that has done more than one installation would know that. Any carrier, that must do hundreds or thousands of those type of installations each year must know that.

So too with the co-lo site. It can hardly be the case that we are the first customer of theirs ever have a circuit delivered by a major carrier. They would have carrier field techs coming in and out on a daily basis.

Yet, the 'demarcation problem' resulted in:

a) the carrier tech, through their provisioning department, telling us that the order has been put on hold and has to be rebooked because of site issues

b) it took a lot of digging and questioning about what the issue was to find out that the co-lo provider had told the field tech he could only run the cable to the co-lo

c) the carrier telling us they could do nothing more until we sorted it out

d) the co-lo provider telling us they had no order for a cross connect, although it was implicit in our original order, and it would take 'three to four weeks if we ordered today to get the cross connect fibre in place'.

I mean, really.

Still, even if it were the case that two carriers, both who according to their marketing are one of the largest four WAN providers in Australia, have never encountered these issues before, I certainly have (on just about every installation actually). Thus within that unhappy MÃ©nage Å trois, one of us, being me, was able to apply quite a bit of hard earned experience in manipulating the sort of lackadaisical, hands in the air, not my problem, wankers that are the architects of the situation. And let me tell you, when you have done that for twenty years, it is a very tedious and boring process to go through 'yet one more time'.

Long story short, the circuit, due today, will be delivered today or 'Monday at the latest', or so I have been promised.

Posted by Steve Waddington at 11:53

Thursday, May 29, 2008

Events of the Week

This week sees the commissioning of the Perth POP. Vina has completed the last of the router configuration, and Raymond and Andrew have updated the database and radius servers respectively. All that remains now is for Telstra to activate BGP and we will be ready to connect our first WA customers to the new POP.

My home connection should be switched over on Friday - it may be the first connection, or, as usually happens, it will probably turn out a few customers are connected first. But either way, Friday should see real production traffic spanning the country on our own WAN network.

Migration to the new ESR's continues. The second 7301 LNS at the Verizon POP was decommissioned last week, and its connections added to the ESR without any problems. Last night as a result of the TW planned outage, all circuits were disconnected from LNS B and E and failed over to the Verizon ESR, which ended up with 22,000 connections at one point.

Of course, there wasn't enough bandwidth on that one TW GE cross-connect, so the ESR was seeing continual VPDN tunnel drops, driving the CPU sky high (as you can see on the MRTG graph). But other than that, and if that were not the case, it would have had no problems handling that many circuits.

The decommissioned LNS will go to Ahad to be used as a dedicated router for metro Ethernet connections.

We had planned to do that earlier, but had to use the router earmarked for that as a replacement for the 7301 that had the failed power supply a few weeks ago.

Incidentally, we had that power supply replaced for just \$300, and we were going to use that router as the new office CPE router. Unfortunately, when Ahad was setting it up today he found that the router now hands on reboot sometimes and shows a few other weird characteristics. It's dead Jim.

Our plan is to move the network to all iBGP. I started to make the changes that would migrate the Powertel and Verizon POP cross-connect on the weekend, however it didn't go as well as I had hoped.

I ran into an issue with propagation of static routes, and found that to ensure correct routing to subnets, some fancy filtering would need to be put in place.

Not that there is anything wrong with using access-lists in an iBGP network - if they can be standardized and don't need to be changed except when the network topology changes.

What I found though was that, to overcome the possibility of incorrect route propagation, each router in the iBGP community would need to have its own unique access-list. And, that each one would have to be updated, not always in the same way, each time a new subnet was added.

Such a thing, while it would probably work, is just not maintainable in the long run. It would be hard to create a standard process to control, and because of that there would almost certainly be errors made down the track, not the mention the time needed to maintain it properly.

So, I scrapped that idea and went back to the drawing board (or the graph paper in this case).

What is needed (which I thought I had but was wrong) is a route policy that basically takes care of itself. That is the goal for any network. We certainly had it when we had just one POP running EIGRP.

There are two things that make the route policy difficult:

- the subnets assigned to ADSL customers
- maintaining all services while changing the underlying route policy and routing protocol

However, after my 'practice run' last weekend, I think I now have the problems solved. The two main problems that held

me back was first that EIGRP has a different default metric to iBGP, and when a subnet route appears in both, it causes a RIB failure, so it appears as if iBGP isn't working, when in fact it would be if EIGRP were switched off.

Second, iBGP flags a route as 'incomplete' if it is just redistributed as a connected route, or there is not a specific route to the subnet. Incomplete routes are added to neighbor route tables only and are not forwarded by those neighbors to other BGP peers. But if a specific route is added to the LNS routers, then that route is also propagated via EIGRP, and so poisons the correct route learned in the rest of the network, which black-holes the subnet on part of our network.

The first way I thought of to overcome that was to add access-list filters on each LNS to prevent route poisoning. Then I realized that a reverse filter would be needed on the core routers, and yet another filter pair for the border-to-core and border-to-LNS peers. Just not workable.

The solution is to use filtering, but specifically for EIGRP rather than iBGP. While complex, it only needs to be put in place until the iBGP process is proved for one, and then for all subnets, then it can be removed, or rather EIGRP will be removed since it will no longer be needed. Essentially what will happen is that through filtering routes will be removed from the EIGRP RIB as they are added to the iBGP RIB. Everything that is running on EIGRP will keep doing so, until it is blocked and then added to iBGP. The end result is that ALL routes will be blocked in EIGRP, which can then be removed totally, as ALL routes will then be carried by iBGP.

The good thing is that we have time at the moment to make sure there are no unexpected surprises from the planned changes. The network is in a less 'stressed' state than it has been for a while, thanks to both the ESR LNS's now in production. That means even if my current planned change iteration doesn't do what I think it will this weekend, or something crops up, there is no urgent pressure that it has to be done. What we will have is a very robust and maintainable route policy when it is done though. Well worth taking a bit of time over.

Posted by Steve Waddington at 07:27

Wednesday, May 28, 2008

Words Fail

I was copied on a summary report this morning from Brendon that showed the total credits that have been approved from one of our carrier suppliers since September 2006. Now we all know invoices with many lines, for many services are going to be a bit wrong - start and end dates for some services might not match up, price changes that were agreed might not have made it through to billing before the invoice run, and a few errors are always bound to creep in in some of the non-automated parts of the billing.

But, at the end of the day it should be all swings and roundabouts. Errors in favour of the supplier should, over time, roughly balance errors in favour of the customer. At least, wouldn't you think?

I mean, business is conducted on the basis of 'good faith', suppliers charge the agreed price and try their best (with the inevitable error here and there) to bill fairly. Customers accept the bill on face value, and pay it on time (or at least Exetel always has) with the expectation that it is mostly right, but if it isn't the errors will a) be fairly minor b) even out and c) can be claimed as credits when/if they are found.

Ha ha ha. Yes indeed, that is exactly how it works, you poor naive soul.

So I read the figure at the bottom of the report and I wondered "Is that reasonable, or what a reasonable person would think reasonable?". Not trusting my own judgment, I walked down to the kitchen and asked Barbara, a very reasonable person, to "take a guess at the credits approved over the last two years" by this particular supplier.

"150 thousand?" she ventured.

"No, higher" I said.

"half a million?"

"Nope"

"Higher?"

"Yes, higher"

"Um, 750 thousand?"

"higher"

"More than seven hundred and fifty thousand dollars? Really? How much then?"

"One point three million" I revealed.

"Wow" she said.

Wow indeed. Having been involved in many parts of the credit claim process over the years, I knew it was in that ballpark. But seeing the one total figure is still quite shocking. And bear in mind that this is the credits approved and actually paid - not the total claimed and/or yet to be paid (which is quite a bit more).

There is no question that only two things have allowed us to stay in business against that sort of practice (perhaps it would even be described in a court as 'mal practice' or 'fraud'). The insistence by John that we automate the bill checking (and the near miraculous feat the dev team did actually building it), and the sheer bloody mindedness on the part of our finance team to pursue the credits we are owed against all the bureaucratic inertia, arse covering denial and soul destroying reconciliation process that a very large company can bring to bear to grind down their much smaller customers. Some of those credits took over nine months to agree. Some that were agreed took over six months to actually appear on an invoice.

One million, three hundred thousand dollars and some change. Shocking isn't the word for it, disgraceful? shameful? criminal?... words just fail.

Posted by Steve Waddington at 10:10

Thursday, May 22, 2008

The strange case of EIGRP flapping - Pt 2

Ok, I think I figured it out. Here is how it could have happened:

Step 1 - One of the routers that was flapping was the edge router we replaced the week before because of a failed power supply. We didn't have the config saved due to a backup server error, and so the config was re-done manually. When that was done, the GE 0/0 interface MTU was left at 1500, and, since there was no evident problem at that point, the omission wasn't picked up.

Step 2 - when the new ESR was put into service, it had no connections, and so the EIGRP packets exchanged with the edge router were not many, and always under 1500 bytes of information.

Step 3 - When the old LNS router was taken out of service to make way for the new ESR LNS, several dozen, or hundreds of circuits would have been reconnecting to the other LNS in a few seconds. The EIGRP information packets would have therefore been large - as big as the MTU allowed in fact - and from the LNS to the edge router would have been 1600 bytes.

Step 4 - the edge router interface however only an MTU of 1500. If the packets weren't fragmenting, and it seems they weren't, then the edge router might be seeing a few HELLO packets, expecting to get a bunch of routes, but then receiving 1600 byte packets, dropping them, and so never getting the updates it was expecting, timing out and eventually resetting the EIGRP session with that neighbor.

Step 5 - in the meantime, the ESR, also with an MTU of 1600, was slowly picking up connections - at around a rate of one or two a second, but never enough to fill up an EIGRP packet above 1500 bytes. So the edge router had no problems with that EIGRP session and it stayed up. The other router on that POP LAN was the second WAN edge router. It also had very few route updates, and so it also did not have a problem with needing to send packets larger than 1500 bytes and its EIGRP session remained stable too.

Step 6 - the reason that no router had a problem with the edge router MTU earlier, is because when it was replaced, after six hours of downtime, there were not many connections on that POP, and when the link was up and they started to build up again, they only did so at one or two a second - not enough to push the EIGRP packet size above 1500 bytes.

Looking back in the router logs, there is evidence of some EIGRP flapping at that time, but with a big disruption it is not uncommon to see that. It settled down after a few cycles and didn't cause any further problem, and so was not further investigated.

That is a fairly long chain to form a conclusion from, but I can't see what else it would be. A way to test would be to change the MTU on the edge router back to 1500 with the expectation that all the EIGRP sessions should remain stable. Then 'force' several hundred disconnections on one LNS, which would reconnect to the other LNS, at which point the EIGRP session between that LNS and the edge router should time out and flap.

I may do that one on some 6am Sunday maintenance window in the future if time permits, because, if it does turn out to be the case, it is information worth knowing. In the meantime, I will give it a label of 'plausible' and since it ain't broken (now), not try to fix it more.

Posted by Steve Waddington at 10:14

Tuesday, May 20, 2008

The strange case of EIGRP flapping

Bringing the ESR into production didn't go quite as well as I thought it had yesterday. Once the old LNS was swapped out, we noticed the traffic pick up again as expected, but after an hour or so, it had not recovered back to the level it should have. It was very strange though, because everything else checked out; there was no problem with connectivity, no errors reported on any interface, circuits were connecting and authenticating correctly, and on the circuits we could test to, there was no apparent packet loss or latency.

Then we noticed that, on the edge router, all the local EIGRP sessions showed an age of 1 hour 20 minutes, except for one LNS router that had an age of only a few seconds. Observing it for a couple of minutes showed RTO for that neighbor incrementing and the session resetting every minute.

If it had been the new ESR that was the neighbor with the problem it would obviously be due to some overlooked configuration problem. However it was one of the existing LNS's, that had operated and peered with the edge router just fine, and for many months. Nothing had been changed on either that LNS router or the edge router.

We looked into it further and first of all determined that there was no connectivity problem between the two routers (why should there be, they sit right on top of each other in the rack and are connected to the same LAN switch), there were no interface errors on either router, no cache flushes or packet loss. What's more, the edge router and the LNS EIGRP neighbor sessions with the two other routers on that LAN showed no problems at all. If say all the sessions to that LNS were flapping, then we could narrow it down to a specific problem with that router. But the other sessions were just fine - it was only the one between the LNS and the edge router that had the problem.

The next troubleshooting step was to check with the last config backup done, the 1am snapshot in this case, against the current running config. It turned out to confirm what we thought was the case - that there had been no config changes on either router.

Usually with EIGRP problems, there is some sort of clue in the log file, like SIA messages, or duplicate IP address messages or something. In this case there was also nothing. The logs on both routers only showed the loss and then acquisition of a neighbor, with no other clue as to why. Very useful.

So we had two routers, with exactly the same config as yesterday when everything was fine, now suddenly having a problem, seemingly, because another router had been added to the network, yet the new router, and all the other routers had no problem at all. Very strange.

Just on the off chance something may have been done to the switch ports configuration, we checked that too - no change also.

We had now moved into the twilight zone in terms of weirdness of the problem. Obvious causes had been eliminated, as had probable and even plausible causes. It was time to check for any difference between the edge and LNS router, and the other working routers on the network.

A few minor differences were turned up, and removed. And made no difference to the problem.

We considered that the old LNS the ESR had replaced might be causing a problem (who knows how, but maybe), since it was still on the network, albeit with different IP addresses and now carrying no traffic, and was seen as an EIGRP neighbor by the other routers. We shut down all the not-in-use interfaces and the EIGRP session. But it wasn't that, it still made no difference.

We then noticed one difference between the edge router and the LNS. The edge router's GE LAN interface had the default MTU, while the LNS's LAN interface had MTU set to 1600. (we commonly set a higher MTU within our network because we often carry L2TP traffic between routers that could have a packet size up to 1535 bytes). However, all the other routers on the network, except the edge router also had an MTU of 1600, and they had no problems.

Never the less, it couldn't hurt to change the edge router MTU to 1600. So we did. And the problem went away immediately and the EIGRP session has been rock solid since.

The summary of the problem was:

- something that had worked with no problems for months, stopped working, and there had been no changes
- except for a new device that had been added, but that device had no problems
- and everything else on the network, with the same configuration as before, had no problem
- the problem was resolved by changing a parameter that should not have had any effect in any event

I seem to recall OSPF is a bit funny about neighbor link MTU sizes. Maybe I had forgotten EIGRP was too? So I spent some hours last evening reviewing the CCO EIGRP white papers, design guides and trouble shooting pages. Lots of information about SIA, nothing about any issue being caused by MTU, at least that I could find. And what's more, it had worked for ages with a mis-matched MTU anyway.

Well, even though we found and fixed the problem, I can't say I'm happy about it, because I really want to know just why it was. We have opened a TAC case to see if they can shed some light on it, but, after the fact the problem is going to be very hard to track down, so I am not hopeful.

Posted by Steve Waddington at 09:57

Monday, May 19, 2008

New ESR in Production

The swap-over of circuits to the new ESR went very smoothly. We decided to take it in stages, as we did with the first one, and just have it replace the function of a single LNS to start off with. That meant we could do a soft circuit swap, rather than needing a few minutes outage.

The first 500 end user circuits on the new box are showing no problems. It bodes well for the full service swap out we will schedule later in the week, or maybe early next week. We have a carrier/supplier switch relocation to get through later this week and it is never a good idea to schedule too many changes too close to each other.

Posted by Steve Waddington at 15:49

Friday, May 16. 2008

2nd ESR Goes in Today

All the components, rack and cabling has come together and we will be taking our other ESR to its new home at our second Sydney POP today.

The performance of the first one we installed last month has been great - it has exceeded my expectations and with now 11,000 active connections is running at an average of just 6% CPU utilization. It is clear that neither unit will have any problem handling the full NSW connection load of ~20,000 each.

We plan to make the circuit switch from the 7301 LNS's to the new ESR early Monday morning. Unlike the last change over, this will be a 'hard' swap of circuits, and will cause a few minutes disconnection for some end users.

We will also take the opportunity to change the intra-POP routing protocol from EIGRP to iBGP. That will be another milestone progressing towards the Foundry NetIron MLX we plan for our network core later in the year.

Posted by Steve Waddington at 09:00

Thursday, May 15, 2008

Meetings

Meetings are like the drop of oil on the gears of a well maintained machine that keep it running smoothly. At the 'one drop' level they are useful to keep the wheels turning, and if they were totally absent, there is a good chance things could go wrong even for the best engineered machine.

More than one drop however, doesn't make any difference to the machine running, but by the same token, doesn't help it run any better.

If you have a choice between having your production engine seizing up, or having perhaps one or two meetings that aren't absolutely necessary, it is undoubtedly prudent to err on the side of too many than too few. Yet it should be within the auspice of every aspiring manager to develop a sense of the 'just right' balance; that does not waste the companies precious time resource, yet keeps the wheels turning in the most efficient way.

Never the less, there should never be a meeting held without a clear agenda and a firm idea of the desired result that directly address some problem. In a notional 60 hour working week, one hour of meetings is probably too much (at least during 'business hours').

Not quite the same topic, but I had another request today from a potential supplier requesting a 'meeting' to 'get to know each other'. I asked for an agenda for such a meeting and received:

The Agenda items for the meeting are:

1. History of both business to get an understanding of what each other does.
2. Focus of each business moving forward - where each of us is putting their attention for maximum market traction.
3. Discuss previous conversations about requirements for Exetel in terms of the following products:
 - co-location in Adelaide and Perth
 - 50-100Mbps backhaul from Perth to Sydney and Adelaide to Sydney
 - 50-100Mbps Internet access in Perth and Adelaide
4. Exetel services that you maybe able to provide xxxxxxx
5. Any other topics of discussion.

So ok, a junior rep, with what looks to me like guidance from a sales manager stuck in the 80's.

My reply was:

"Dear Sir,

I can give you the information you are looking for, interleaved below:

> >
> > 1. History of both business to get an understanding of what each other does.
> >

The Exetel web site = <http://www.exetel.com.au>, has a complete history of the company. I can also get any information I would want from your web site.

> > 2. Focus of each business moving forward - where each of us is putting their attention for maximum market traction.

The Exetel web site has that information

- >>
>> 3. Discuss previous conversations about requirements for Exetel in terms of the following products:
>>
>> - co-location in Adelaide and Perth
>> - 50-100Mbps backhaul from Perth to Sydney and Adelaide to Sydney
>> - 50-100Mbps Internet access in Perth and Adelaide

We currently have two year contracts with another company for those services.

- >>
>> 4. Exetel services that you maybe able to provide xxxxx

Can you please let me know what you are looking for and I will put you in touch with an account manager who can quote you on the price.

- >>
>> 5. Any other topics of discussion.
>>
>> What time Thursday can you allocate say 1.5 hours to discuss?

If you still wish to meet, I think we can cover anything else in about 15 minutes. I will be in Sydney next week and mornings are the best time for me.

Because, from their web site, it did look like they might have services that could be of interest to us. And the sales people of a good technical company are not necessarily an indication of the capabilities of that company.

I received a short response that seemed a little petulant to me, but no further requests for a 'meeting'. I think I will count that as an efficient conclusion.

Posted by Steve Waddington at 15:46

Tuesday, May 13, 2008

Dead? What do you mean dead?

... is usually my first response when someone says 'the router is dead'. There are lots of ways a router can be or appear 'dead', interfaces can lock up, it can be CPU bound, or run out of memory. All those things will stop it routing traffic, make it unreachable, and is generally what people mean when they say it is dead. And unless there has been some physical fault with the router flash, a power off and power on will restore it to working condition - a wait of 2 minutes and 33 seconds for a 7301 class router.

Then there is the other 'dead' - no lights, no fans, yes the power is plugged in. but there is no difference between the power switch being on and the power switch being off. It is the worst possible scenario. There is no hope of returning the router to service this side of factory refurbishment. The only thing to do is to get a new router from the spares stock and configure it with the failed routers profile.

Because everyone has \$40,000 routers sitting around as spares, right?

Well, the worst kind of dead is what happened to an edge router at our second Sydney POP at 11:30pm on Friday night. Being the router that terminates the WAN circuit for the two POP's it meant that a) we had lost 1/3 of our external Internet bandwidth, b) VoIP services were down c) customers connected via that POP - about 7,000, had no external access. A fairly major problem in other words.

One slight upside was that customers connected at the second POP could restore their access by simply disconnecting and reconnecting their ADSL circuit (power cycling their modem being the most common method). They would then automatically be reconnected to the man POP with full access to the Internet and all other services.

At 11:30pm when all the alarm bells went off, it first of all looked like a WAN link failure between the two POP's. The duty engineers at the POP were again very helpful, within about 15 minutes one of them was in front of our rack and talking to Vinna on the phone telling her what he could see. His report was that the WAN switch from Pipe (who provide the POP cross-connect circuit) showed an active LAN light but the WAN light was off.

The next step was to contact the Pipe duty engineer, who promptly confirmed the condition and ascribed it to a faulty cable (a similar failure we had had recently on that same link was due to cabling). The POP engineer was right on the job and ran a new cable for us, but no joy.

Maybe it was because it was late at night, but we had missed a very important piece of information. The POP engineer had reported that the LAN light was on, but the WAN light was off. Pipe had confirmed that yes, the LAN segment was active and the WAN was down. Both were correctly telling us what the problem was - which was the WAN connection - but it was the WAN connection from Pipe to us, which it to us a LAN cable between our router and the Pipe switch. Pipe's LAN connection is therefore our WAN connection. So what both engineers were telling us was that their LAN connection was fine, and it was the connection between our router and their switch that was the problem.

These things happen I guess. But there was 45 minutes of downtime looking at the wrong thing we could have avoided.

Ok, the penny had dropped and we now had the POP engineer back in front of our rack looking at the cable from the Pipe switch to our edge router. 'The router is dead' he said to Vinna who reported it to me, leading back to the first three paragraphs.

Fortunately, we have always planned 1 to n redundancy into the network, so we did in fact have a spare router we could replace the dead one with. It was just at a different POP. Given that it was now getting near to 2am, non p2p traffic had eased to the extent that it could be easily carried on the remaining available bandwidth. Connections to our main POP LNS's also showed that anyone who had lost access had by now rebooted their modem and were reconnected. That only left VoIP services, which are at a low from midnight to 10am.

It was midnight for me in Perth, and I was booked on the 6am flight to Sydney the next morning. Normally I would have just canceled the flight and provided remote backup to Vinna while she swapped out the router. However, since our replacement router hadn't been tested, there was some chance, with transit and all, that it too could have a problem. My flight to Sydney was via Melbourne, where we also have a spare router. So what I decided to do was go ahead with the

flight, meanwhile Vinna would get a few hours sleep and pick up our Sydney spare router around 5am. If everything went well, she would have it replaced by 8am - and acceptable emergency replacement window. In the meantime, I would retrieve the failed router config from daily backups and send it to her ready to download to the replacement router when she got on site.

Then, in the worst case of the replacement router not working, I would pick up the other spare in Melbourne and bring it to Sydney.

I really didn't like the idea of being on a plane with such a major problem unresolved, but it seemed like to best, safest, course of action.

So while Vinna got some shut-eye, I logged into the backup server to get the router config from the day before. Zero byte file. hmmm. The day before that. Zero bytes. A week ago, still zero. Grep for anything larger than zero. No joy. Every other router config backup was fine, just not the one needed. It turned out later to be a hosts.allow permission problem, easily fix, but that was no help at the time.

I sent Vinna an email with the bad news, but encouragement that at least an edge router config was straight forward. In fact all that was really needed was to put in the right interface addresses and enable eigrp and everything would take care of itself. A few other things occurred to me that might be helpful, so I did what I could and then went to pack.

I called Vinna from Perth airport at about 5am (7am Sydney time). She had got to the POP an hour and a half ago, but had only just got access to the co-lo room. Whatever monkeys those guys hire now were not able to activate the remote access release to let her in and had stuffed around for over an hour before they had figured it out.

I had hoped she would have had the replacement router physically in place by now, and I could help her out before my plane took off if any unexpected problems had cropped up. That was still at least an hour away, my plane was boarding in 15 minutes and my fingernails were down to the second knuckle.

It was a very, very unpleasant flight to Melbourne. Not because the seats are so bad on those old 737's Qantas think are suitable for domestic use. Not because as the plane was taxiing, it turned around and went back to the boarding gate because of an 'instrument problem' (the flight engineer came on board, pressed ctrl-alt-delete or whatever and the plane taxied out again) a mere 40 minute delay. It was the four hour 'communications black hole of worry' while my mind worked on all the possible things that could go wrong and the ever more depressing situation I might find on arrival in Melbourne. I think I mapped out about fifteen alternative network configurations we could use if the worst of the worst case possibilities I was imagining ever eventuated.

As soon as I could switch my phone on at arrival in Melbourne it lit up like a Christmas tree with sms messages arriving. Most were from the alarm system telling me what I already knew. The ray of hope, the POP WAN link was below threshold - a promising message because when there is no data on that link, the threshold monitor says nothing, only when traffic is flowing can it report a lower than expected state. Usually when a link comes back up, a few low threshold messages are sent while the traffic picks up.

Then a message from Vinna, what had gone wrong? What terrible, unforeseen and unrepairable event had taken place while I was in transit?

The message read: "edge router returned to service before 9am"

By this time my laptop had booted and connected to the airport wifi and I could see the replacement router carrying the normal time of day traffic, and all services restored. What a relief. Of course I knew all along it would be done on time with no further problems.

As I remarked to John a little later, that was 12 hours I really could have done without.

Posted by Steve Waddington at 14:57

Thursday, May 1, 2008

ESR Results

The ESR is performing very well. Replacing one 7301 running at about 60% CPU with 5,000 connections, the ESR has been doing the same job at 2-3% CPU utilisation.

Today Vinna switched over connections from a second 7301 LNS, adding another 5,000 connections and bringing the total traffic to 800/600 Mbps (you can probably tell they are ADSL2 connections). We also added a route map for one IP address to direct it to its own egress path - usually a sure way to push a router's CPU to an unsustainably high value. The ESR however sucked it up and went to almost 6% CPU utilisation.

We have two more LNS's to cut across, and the early indicators are that after that the ESR will still be able to handle three times more connections.

Posted by Steve Waddington at 20:05

Saturday, April 26. 2008

Recovering lost time

Sometimes, in fact most times, most people will accept delays and hold-ups that are beyond their control and just live with the consequences. I mean circumstances like when the train you are going to work on breaks down, making you miss a 9am meeting, or your plane is delayed causing some other planned event to not happen. It isn't as if there is much that can be done about it by the individual affected. So, while it might be frustrating, there is no point getting upset about it, these things just happen.

Then there are times when, no, those slings and arrows will not be suffered. What is planned to happen will happen despite them. Will will not surrender to fate, this time fate will be shaped to will. I would think anyone that has sailed outside a protected harbour will know what I mean, and understands exactly why wind and sea are used as such strong metaphors for fate and fortune. You accept not where the wind and waves would drive you, but hold the tiller firm and steer your course into the teeth of all that comes at you regardless.

Sometimes, that safe harbour of 'I would have completed it but for...' is just not an option.

So on Tuesday, receiving confirmation that the cross-connect cables to our new rack had finally been completed, I had to decide if we could:

- bring the ESR into service that day, so that
- we could free up at least one of the 7301 LNS routers, so that
- It could be take to Melbourne on Wednesday, so that
- work could be completed in an early morning maintenance window on Thursday morning, so that
- I can get the last available flight back to Perth on Thursday, so that
- I can take my boys to the ANZAC day dawn service on Friday

And, it all had to be done with minimal risk of error and no noticeable disruption to service.

Our careful plan, drawn up and refined during the long wait for the cabling to be completed, called for the cutover of traffic from two LNS routers to the ESR to be done, with a little disruption, but not too much, during an early morning maintenance window. At first glance, it just did not seem possible that that could be accomplished in the time available. However, it then occurred to me that if we cut over just one LNS, essentially duplicating its function on the ESR, and timed it just right, we would have the router needed for Melbourne, and not cause a noticeable impact on service.

If it were just the Melbourne POP upgrade itself, I probably would not have decided to go ahead with such a tight schedule. But since the project was already a month late, it had really impacted on the whole national network deployment plan, affecting the planned POPs in Perth and Adelaide, the upgrade to the 2nd Sydney POP, plans for Sri Lanka and a host of other minor jobs. The cascading delays it caused had to stop. Sipping my morning coffee on Tuesday I could see a clear path through the morass that would achieve all the objectives we wanted. Tight, yes. Required a few breaks our way, yes. But do-able.

Step one was of course to connect up the cables and bring the ESR on-line. It was also the first of many possible stumbling blocks, as everything else was dependent on the cross-connect circuits actual working.

We were at the POP by 10:30am, plugged the cables in, link lights all went green. No cell flow. Check the switch ports - up, check the router ports - up/up, all looks good. No cell flow. Plug in the laptop, plenty of cell flow. Plug back in the router, nothing.

Fiddle with the router and switch port settings. There aren't a lot of controls on a GE interface, basically it can be auto-negotiation or no auto-negotiate. Neither made any difference, still no cell flow.

The only possibility was that it could be the GBIC's we were using in the ESR might not be compatible with the switch port interfaces. A call to our supplier confirmed that yes, they should work. But since they weren't we asked if he had others. Yes, he had exactly the three we needed of the alternate model in stock, he would ship them tonight and we should have them tomorrow.

This is where five years of doing business with someone pays off. We had to have the GBIC's today, just about within the hour, if the schedule was to be met. Hamid came through for us, grabbed them from stock, and drove them to us himself. By noon we had the replacement GBIC's. And they worked!

Ok, now a slightly risky part. We duplicated the config of the 7301 LNS on the ESR, which meant duplicating loopback and supplier interface addresses. We had done this before by accident during an earlier change over some time ago and found that, while the routers complained about the duplicate address, it made no functional difference since the interfaces were separated by the core network.

The tricky part was to put the ESR interface on the same VLAN as the 7301, shut the 7301 interface, no shut the ESR and change the core router static route to point the traffic from the 7301 to the ESR instead. Not exactly simultaneously, but within about a second of each other. If we did that for one VPDN circuit first and it worked, it would work for the others.

Our test case VPDN is one that currently has only two staff accounts connected. We made the change while pinging the end user modems, and it went perfectly. One packet drop on one, and two on the other - likely not even a second of interruption.

Next we swung over a more substantial VPDN circuit, this time with about 20 end users connected. Again success with no noticeable interruption. We then did some more, bring the total connection on the ESR up to 270. Then it was time to leave it for a while to see a) if the ESR was performing as expected and b) if there would be any problems reported by anyone connected to the ESR.

We gave it an hour and there were no problems reported or evident. Another 400 circuits were cut across, another hour, and still no problems. We then began gradually migrating the rest of the circuits off the 7301 while monitoring the ESR for any sign of problem. At this stage, if there was a problem, it was a simple and painless roll-back to the 7301. By 6pm all the circuits has been cut across from the 7301 and the ESR was performing as well as we had hoped.

The next step of the plan was to let the ESR run overnight with the 7301 in place, and if there were no problems, I would remove the 7301 in the morning on the way to the airport to catch the 8:30am flight to Melbourne. I went of home to have some dinner and pack.

Since the circuit in my Sydney flat was one of the ones migrated, I was testing it out that evening to see if I, as a user could detect any problems. At first it seemed all ok, but then I noticed some web sites were slow to load. No sooner had I started investigating that I received an SMS alarm about low threshold traffic. 'Oh no, the ESR has struck some unexpected problem and we will have to roll back' was my first thought. However further investigation a few minutes later showed that the problem seemed to be packet loss between the two Sydney POP's. By that time plenty of other customers, not connected via the ESR were also reporting problems. It didn't fully rule out the ESR as the cause, there could be some weird routing or route flap issue with it affecting the whole Sydney network. The evidence did seem to point to something else though.

As it turned out, it was a supplier switch causing the problem. The supplier rebooted the switch at about 9pm and the problem went away. By 11pm there had been no recurrence of the problem, and I was sure the ESR was not to blame. Never the less, I no longer had the confidence to remove the 7301 at 6:30am as I had planned. Once that router was gone, there was no possibility of rolling back if there was a problem with the ESR. I needed an alternative plan.

My about midnight I had mapped out what could be done. I needed to confer with my colleagues at the office in the morning. If we all could see no further problems by 11am I would remove the 7301. We could also in the meantime configure another 7301 as a backup to give us nearly the same roll-back path as if the removed 7301 was still there. Then, if all was well, I would take the 2pm flight to Melbourne, be at the POP by 4pm, install the router, make the 5am window and be on the 8:30am flight on Thursday to Perth (only flight available to Perth from Melbourne on that day).

There were no more problems. I made it to Sydney airport at 1:15pm with the 7301 router in an anti-static back in my travel case. The 2pm flight to Melbourne was due to board at 1:30pm, but the board showed it delayed by 20 minutes. No problem, a few minutes later didn't matter.

1:50 came and went. At 2pm, the scheduled departure time, Qantas announced the flight had been canceled and 'all passengers please rebook at the service desk'. Ah, dammit, my plan was shot.

Because I was standing at the back of the boarding gate lounge, I was one of the first to the service desk. I was re-booked on the 2:30pm flight. I moved the schedule forward in my mind half an hour, and yes, it could still be done. If

there had ever been room for error, there was absolutely none now.

The 2:30 flight was only ten minutes late taking off. There really was nothing I could do about that, so like Jake and Elwood in the elevator on their way to the County Assessors office, I just listened to the music. That music must have triggered a few synapses, because somewhere over Canberra I figured a way of installing the new 7301 in Melbourne with zero risk of disruption to services there. That meant I could complete the work in one go, without the need to complete it at 5am - something I wasn't looking forward to since I was beginning to feel a little fatigued for some reason.

I made it to the Melbourne POP at about 4:50pm and started work. The router was no problem to physically install, except for some poorly thought out cable placement from previous work. By 6:30pm it was plugged in and powered and a check back to Sydney confirmed there were no problems and everything was ok.

The switch reconfiguration similarly went very smoothly. That is the beauty of dot1q LAN's - they provide a great deal of topology flexibility and allow changes to be done without interruption that would otherwise require significant down time. The 7301 was similarly reconfigured with the planned code. I checked out the modifications I had thought of on the plane and there was nothing I could see that would cause a problem.

The software work was done by nine. I ran another series of double-checks and it all looked good, still. At 9:30 I made the first change to bring the router into production, and completed the last change at about 10pm. It all went swimmingly. If there was even a second of disruption, it wasn't evident. Whew.

Now, just monitor it all until, say 1am, and I think we are done.

Since that can be done just as easily from the Hotel, I went and checked in and had a late dinner.

There were no problems. It had all worked. We had outlived the day and come safe home.

Posted by Steve Waddington at 12:54

Wednesday, April 23, 2008

Up, up and away

At long last our first ESR 'super router' is in production. We were able to migrate the first few test connections yesterday, and seeing no faults or issues, progressively migrated more during the day.

It was a great result by our small team working under increasing pressure due to the other projects the ESR migration was holding up.

We received the cross-connect completion notice late on Monday but were still unsuccessful bringing up the link to the new unit. On site we were able to see all interfaces reporting 'up/up' but no traffic was flowing at L2 or L3. We seemed to have another one of those frustrating mysteries where the equipment reports no problem, but, like Catweazles magic, nothing works.

Connecting a laptop interface to a switch across the cross-connects showed there was no cable problem (for a change). The only possible cause I could see was an incompatibility between the GBIC's and the media converters we were using.

The GBIC's we had installed were dual use, which meant they could be used for both Single Mode and Multi Mode cable. Or not, as it turned out. We made a call to our supplier who confirmed that they should work, but also had some multi-mode only units in stock. He was kind enough to pack them up and deliver them himself - a favour we greatly appreciated, thanks Hamid.

Ninety minutes later we had the GBIC's swapped out and a 100% success rate on a million ping packet test between the ESR and core network.

IBGP was brought up, and with no little trepidation Vinna pushed in to code to switch over the first group of customer circuits.

"VPDN tunnels are forming" I said

"But show caller summary shows no users connected" said Vinna, followed by "oh, radius is being rejected"

A quick dash around to Andrews desk and whatever it was the radius server didn't like was fixed. The first end user connections established and instant later, and I confirmed their IP's had global reach. The ESR was on-line and carrying production traffic with no problems evident. Success at last, and a nice piece of work for the day.

Posted by Steve Waddington at 05:33

Tuesday, April 22, 2008

Spam can not be beaten

... as long as there are moron sysadmins out there who can't even spell RFC, let alone actually having ever read one.

And I have been reminded just how many moron sysadmins there really are. An underlying principle of the Internet is 'be conservative in what you do, be liberal in what you accept from others'. It may seem like one of those throw away lines, but Postel was a very smart fellow, well respected, and used to dealing with a peer group of similar, respected and respectful, engineers. So when he made that statement, way back in RFC 793, the people reading it at the time gave due consideration to the thought process that would have caused it to be made.

It is only as a direct result of such due consideration to the profound concepts that created the first Internet protocols that the Internet is as incredibly robust as it is today. Allowing it to cope, admirably, with the zillion fold increase in connectivity and traffic, with essentially the same initial design, as it began with. It would be the same as if your brand new XY Falcon in 1971 was now an intergalactic space ship - or think of some other amazing analogy if you are not a Ford fan.

Why am I picking on sysadmins today? I'm not, I am picking on moron sysadmins, those idiot savants that think that just because they have conned an employer into giving them a job looking after servers (often with an attendant MSCE or some other equally irrelevant certificate), they are somehow fit to be responsible for connecting those servers to the Internet. And in doing so, through their ignorance blithely expose the Internet and their own company networks to the worst forms of abuse.

The particular manifestation that has raised my ire is this; for the last 24 hours my email inbox has been bombarded by thousands of 'delivery fail' notices from email domains around the world. Obviously, some spammer has used by email address as the 'From' address for spam email they have sent out. The result is, every poorly run server, managed by said moron sysadmins, has been replying to the from address telling me that it can't accept the email because either the recipient doesn't exist, or because it thinks the message is spam. Spam eh, no sh*t Sherlock. And you think the person who sent the spam used their own email address? Thanks for letting me know that, you total d***head.

Hence 'the worst form of abuse'; where the effects of the original spam are compounded, by being resent, BUT AS LEGITIMATE EMAIL, to a completely innocent bystander. Now if I just happened to have my email account on a server run by another moron sysadmin, then what would happen is the messages may well be bounced again for a similar reason, returned to the idiot systems that sent them, bounced again, again and again in a never ending feedback loop.

It doesn't grow out of control today, because there are plenty of good sysadmins, and well run systems, that do the right thing and silently drop these sort of messages, short circuiting the feedback loop. Those systems are the shock absorbers of the Internet. And while there are plenty of them in operation on the Internet, such loops die off after two or three iterations at most. That may remain the case forever, or it may not, but let's hope it does.

The point is though, that it doesn't matter how many well run systems there are, as long as there are morons maintaining idiot policies, no amount of anti spam controls can ever work.

Here's an idea - license the drivers of servers. For a sysadmin to operate a server on the public Internet, they need a public operators license. It works globally for cars, it worked globally for ham radio operators (if anyone remembers). Both are devices that can cause a lot of problems for others if an internationally understood code of operation is not followed. We have the code of operation in the RFC's, demonstrably one of the most robust sets of codes ever created in history. So robust that they have withstood the morons let loose on it for a long time, but to the very great annoyance of many.

I can't, and I doubt any one can say, how much better it would be if only licensed sysadmins were allowed to run servers on the Internet. History however shows us that modern medicine, performed by licensed practitioners, has been quite a big improvement over witchdoctors and voodoo.

Posted by Steve Waddington at 06:36

Monday, April 21. 2008

Surprise Surprise

The long awaited cross connect, 70 metres distance, ordered two months ago, was not installed last Wednesday. We were given a revised date of Friday. And with no completion notice, and our account manager mysteriously not contactable (on a Friday afternoon), Vinna went to the co-lo to see what could be done.

Well, the cables ends were physically there at each rack, but no continuity between them. Obviously the jumpering hadn't been done on the patch panel. Under the 'rules' of the co-lo, only the authorised cable contractor is allowed to access the patch panel. But just having a look can't hurt (and is it our fault if the wires accidentally patch themselves as a result?). However it was impossible to tell where our patch blocks might be, or even if the cabling to the panel was complete. So we had to leave it.

Rather than further unreturned calls, I asked Vinna to go up to reception (since the co-lo and the main office are in the same building), and plant herself there until she could get to talk to someone - anyone, about what was going on. She managed to catch our account manager on his way out (4:30pm). He explained there was nothing he could do since, apparently, all cabling work was out sourced, there was no one in-house who could do it, and the cable contractors had left for the day.

I guess it would be just too bad if there was a cable break on a main trunk circuit. It would just have to wait until the 'cable contractors' got back on Monday. How can any company be taken seriously as a carrier with an attitude like that? I think a new low in telco standards has been reached.

I emailed our account manager this morning. I wasn't expecting anything until my third or forth follow up call, but just received this:

"I was just speaking to Vinna about this one. I have received an update of COB, Tuesday 22 April 2008, however have also spoken to the group concerned to see if it can be completed today. The reason it was not working on Friday was that the patch panels and tie cables had been completed, however not the jumpering to the ID"

So there you go.

It is a very, very painful and expensive exercise to move operations from one data centre to another. Fortunately, my reluctance to contemplate such a logistical horror was overruled by the foresight of the other directors some time ago, so we have a second co-lo we can switch to. Unfortunately however, even with a second location, such a change is still a six to nine month process. Never the less, it is clear it has to be done, so that is what we will do.

Posted by Steve Waddington at 11:21

Friday, April 18. 2008

Tunnel at the end of the light

Following, as one does, the current fiasco about broadband rollouts makes me wonder if we are not coming to the end of what in years to come will be termed a Golden Age of Internet in this country (not that it is, just that it will look that way compared to what is to come).

I mean what chance is there that a government that is completely stymied by the massive complexity of a plastic bag, is going to get broadband right?

But hey, come on, lets be fair. Plastic bags are hard in politics. People completely understand plastic bags, there is no mystery, no technology gobbledygook, no one imagines scary pedophiles and porn pushers beneath every pile. No one thinks Australia needs a Plastic Bag Superhighway be internationally competitive. And, no plastic bag manufacturer donated \$10 million to any party at the last election, at least as far as I know.

So it is a really tough decision for a politician to make. Every voter will, possibly, suffer some very minor inconvenience by banning them, which can not ever be blamed on the previous government (never mind the massive benefit it will have to the planet, it seems the current Federal minister has lost confidence that stage gyrations and cutting political lyrics are good enough to sell policies any more). The Australian public has an arguably perfect understanding of plastic bags, they can't be fooled, and will know exactly where to place the blame for their supposed annoyance at any decision. Scary, scary stuff.

Broadband however, now who understand that, apart from a few loony ISP's and a bunch of geeky nerds. A political godsend for those 'tough decisions' that 'have to be made'.

Posted by Steve Waddington at 11:46

Thursday, April 17, 2008

Light at the end of the tunnel...

...or just a lighter tunnel?

After two months of delay it looks like the cross connect cables to a new rack we ordered in February will finally be completed tomorrow. It is only 70 metres between the two racks. Completion tomorrow will mean the cable has been provisioned at a rate of just over one metre per day. I am trying to imagine something that moves that slow so I can draw an amusing analogy, but I can't (since apparently global warming has sped up glacial flow to faster than that now).

And I hope I haven't now jinxed it by saying that. Fingers crossed.

Completion of the cross-connect will mean we will be able to commence the deployment of our first Cisco ESR into production. One (albeit very small) benefit of the long delay is that we have had plenty of time to configure the box, and test it, and then test it again just for good measure.

I think I deployed the very first of that model to be shipped to Australia, back in 2002. Or if it wasn't the first, we took delivery before any others had made it into production. Well, there have been three generations of processor upgrades since then, and big advancements in line cards (the first models came with your choice of OC3 or OC12 interfaces, except the OC12 cards hadn't rolled out of the factory yet). Whereas the early models, from memory, were pushed delivering 16,000 connections, I fully expect plenty of headroom to be left handling our (modest) 50,000 connections across the two we will deploy.

Particularly since we have over the years, moved all the 'fancy' processing to specialized servers at our Internet edge. Which will leave the ESR's to do nothing but what they do best - terminate L2 DSL connections.

If all goes well, we may have the first trial connections cut over late Friday or on the weekend. As usual, our staff ADSL circuits will be the guinea pigs for that. There is a good chance we will be able to make the circuit cutover seamless, but certainly it should be no more than an early morning few seconds disconnect/reconnect.

That is; assuming the cross-connect circuits are completed.

Posted by Steve Waddington at 16:26

Wednesday, March 26, 2008

The 'Work from Home' Acid test

I've been thinking for a while about how to manage an implement 'work from home' policy (about nine years in fact). Just what are the nuts and bolts of how such a thing should work? I drafted a fairly comprehensive white paper about it last year, in an effort to work out just what it is a company, a manager and the employee needs to understand and agree if it were to happen.

So what is the 'fear' of having a staff work from home policy in Australia? There is no longer any problem with the technology, the full kit for a home office, including IP phone, router and computer, is less than \$3,000. Communications certainly isn't a problem any more, certainly not anywhere in urban Australia, nor in Colombo either.

The two main negatives I see are:

* I don't know if you have noticed, but Australia is a very nice place to live. We all have XBoxes, Foxtel, live near shopping malls and lots and lots of other lifestyle distractions. It might not be the intention of anyone to 'bludge' working from home, just as it isn't the intention of someone to become a crack addict, for those that fall prey. Would the temptation be just too strong?

* Maybe some people are just naturals when it comes to managing people. But I know I, and quite possibly most people, are not. It is tough to do properly, even face to face in the same office. Mistakes are easy to make, and can be hard to correct. The propensity for technical people to not be the most social of people also needs to be taken into account. So is the manager/supervisor, and/or the company management procedures in fact good enough for those inevitable issues to be dealt with just via email and phone?

In the meantime of course, we have run a successful 'remote worker' policy for our two staff in Sri Lanka for the last two years. It took longer than we expected to achieve that success, but I guess, in reality, that was only to be expected. Because, apart from the guys in Colombo having to learn fairly abstract concepts about Australian communications, we also had to learn the techniques to support and supervise people in that situation.

There is also my own situation, where I have worked at the Sydney office for two weeks each month, and the rest of the time from home in Perth for the last four years.

Both of those have worked out ok because a) the Sri Lanka staff are very motivated to continue to work for a well paying, O/S company and b) Being a shareholder, and director with personal guarantees, whose family welfare and income is fully dependent on the success of the business, for some reason keeps me fairly focused.

I think all the pro/con arguments can distill down to two points:

- 1) From the individuals point of view, all negatives, perceived or real, can be overcome where there is genuine motivation
- 2) From the company point of view, work from home must mean increased productivity

Certainly in relation to increased productivity, the quid pro quo of regaining two to three hours of commuting time, the comfort of the home environment, and the inevitable personal interruptions that will take place during the day, must be offset by the an extended, and more efficient, working day.

And so, I thought of the 'Acid Test'; which can be applied to the situation of:

- someone wants to work from home
- another person needs to be hired in the area of the person that works from home

Instead of hiring another person, let the the person who wants to work from home do that, on the basis that the company wont need to hire another person.

It seems like a situation that benefits everyone. (except for the person that wasn't hired of course)

Posted by Steve Waddington at 12:55

Tuesday, March 25. 2008

Support coefficient

One of the things we will need to do as we increase our number of customers is improve the support staff-to-customer coefficient. Because of we double our customer base, the last thing we would want to do is double our support staff numbers - a road to ultimate disaster. Because of course, if we did that, we would then need more supervisors, who would need managers, which is the beginning of a layer of middle management, who would have to report to senior managers etc, etc. And however the opinions and management styles of all those supervisors and managers differed, the one opinion that they would all share is that they NEED MORE STAFF.

In the relative blink of an eye, instead of 60,000 customers and 15 support people, we would have 90,000 customers and be opening our third call centre to house our staff of 300 or more.

Well, that is what every other ISP has done at any rate. So even if we had no idea of what to do ourselves, the one thing that we can be absolutely sure of, is not to head down that path.

For the first three years of Exetel's operation I maintained a coefficient graph that tracked the number of customers in relation to the number of support tickets. Using a formula that gave extra weight to new customers, decreasing the weighting in the second month, it it provided a good guide towards various customer support initiatives we put in place.

The way it worked was any time the coefficient was below 1.2 indicated that the ratio was improving (ie we needed less staff time per customer), and over that meant that we needed more. So we were able to gauge month on month how changes to support, introduction of new products or even changes to the documenting supplied to new customers impacted our operating efficiency. And therefore, when the coefficient rose, actions could be put in place, or decisions changed, to swing it back to a downward trend.

By August 2007 however, due to on-going automation that created tickets in the helpdesk for every issue, the gross ticket count was no longer a meaningful input. The reporting base had been broadened over all anyway, as it had to be for the broadened product base, and it was of greater value to track each product type rather than support over all.

Never the less, it is the same principle we have always used, where a scaling factor is 1:1 is a failure. On going automation must be put in place to ensure that the cost of overheads and infrastructure increases at a smaller rate in relation to the customer base.

GURUS is one, in fact the primary way, that will be done as far as the management and reporting structure of Exetel goes over the next 12 months. However support remains the biggest 'cost centre' and closely coupled to GURUS reporting is putting in place better tools to speed up fault resolution and give customers better and faster access to fixing problems.

The first iteration of an ADSL unified fault analysis tool is about 80% complete. When done it will give our support people a complete 'picture' of any customers current status, as our various systems see it, as well as the results of basic diagnostic tests gathered from both our own, and some supplier, interface portals. That alone should save us the first 30 minutes any helpdesk engineer has to spend gathering data manually for many problem types.

Then the next step after that is really exciting. In fact I think it will completely revolutionize broadband support in this country.

But.... one step at a time.

Posted by Steve Waddington at 14:12

Monday, March 24, 2008

WiMax is No Good, Apparently

This article in CommsDay caught my eye. Wow. That is one in the eye for future wireless plans, maybe. But look at what is cited as the cause of the 'failure'.

1. Non line of site performance non existent beyond 2km. Hardly surprising, we aren't talking AM radio here.
2. 1000ms latency. That is the bane of all wireless data comms, has been for years. It has always baffled my just why that is, I am sure a radio engineer could explain it. Maybe it has something to do with signal lock and frequency hopping.
3. Poor indoor signals, base station coverage reduced to 400M.

The result of that was they were unable to achieve the 10% market takeup of their target 55,000 user base.

Which seems to me to be where the real failure is. I don't know who these Buzz Broadband/Queenslander.net people are, not that that means much, but I have never seem them sited in any 'top 10', '20', '30' or even '50' reckoning of Australian ISP's. Which would put them in the bracket of 10,000 users at a stretch, or more likely less than 5,000. Again, I don't know, there is no information I could find on their web site about that, however it does say that "There has been 3612 visitors since February 1st 2006", so draw your own conclusion.

It seems very, very unlikely that the quoted 55,000 user target base is from existing users. Looking at the coverage list on their web site, I gather is possibly the total number of households in the areas they have coverage. And they were expecting to sign up 10% of those? With plans that are more expensive (by about 50%) than the current Optus/Telstra give-away pricing? Well it is clear to me at least where that business plan failed. It's not the first time, and certainly wont be the last that the technology is used as the scapegoat and blamed for the non-realization of a financial dream.

Anyway, I've become sidetracked. It is really worth looking at the criticisms leveled at WiMax to see just to what extent is does fail to deliver on its promise.

I think where many people go wrong is to view WiMax as a competitor of Next-G or other mobile phone based data services. It is certainly attractive to a small operator to do that because of the fairly wide coverage of a single base station. For example, North Sydney could be covered with WiMax using maybe three base stations, or two for about 75%. Much cheaper than the zillion mobile cell antennas that festoon every second building.

A small number of base stations will work for a relatively small number of users. But many of users who all want high bandwidth are going to mandate many base stations and antennas. And deployed in a pretty smart way to avoid spectrum contention, and latency and jitter as a result of frequency hopping. In other words, all the things that have been eliminated from mobile phone networks over the last 20 years.

Sprint's shelving of their proposed \$5bn roll-out of WiMax as a mobile application is probably as good an indicator as any that it just isn't 'there yet' for the 3G/4G market.

However, look at WiMax a different way.

Wikipedia lists the applications of WiMax as:

- * Connecting Wi-Fi hotspots with each other and to other parts of the Internet.
- * Providing a wireless alternative to cable and DSL for last mile broadband access.
- * Providing high-speed data and telecommunications services.
- * Providing a diverse source of Internet connectivity as part of a business continuity plan. That is, if a business has a fixed and a wireless Internet connection, especially from unrelated providers, they are unlikely to be affected by the same service outage.
- * Providing nomadic connectivity.

Our current rate for a 10Mbps Metro Ethernet connection is \$1,200 per month on a 12 month contract, per customer. Of

which, a very large proportion of the \$24,000 paid over the 12 months goes to the carrier. While a WiMax set up for us to reach six to eight customers in North Sydney might be in the order of \$60,000, for the hardware, directional antennas and installation.

Not 'super cheap', not competing with a mass consumer market product. But a good cost saving for a specific purpose in its specific market.

In those circumstances there would be:

- No spectrum contention or frequency hopping, and so low latency
- Fixed, high gain directional antennas, guaranteeing bandwidth and stability
- and eliminating the 'indoor signal' problem

The only thing it doesn't seem to be that good for (yet) is nomadic connectivity. But then how many businesses really work from coffee shop to coffee shop all day? How many more would think that 6, 8, 10 or 20Mbps Internet access for their office, at around half the current price, would be just fine?

I think that is the market.

Posted by Steve Waddington at 09:52

Thursday, March 20, 2008

Over Easter

I pushed hard for it, but it doesn't look like all the bits and pieces will come together before Sunday to do the LNS migration I had planned. Just too much 'can't be bothered' attitude from one supplier in particular. Three of us have spent the last week trying to get information and confirmations - very hard to do when phone calls are never returned and emails are ignored. 'It's like trying to pin down a ghost' as Brendon remarked in exasperation yesterday.

There are really only two times a year when major infrastructure changes can be planned - Christmas/New Year and Easter. I mean, of course they can be planned any time, and done any time. But to minimise impact on customers, and on sales, work has to be done at a time that causes the least inconvenience for the most people. Practically, that is between 6am and 8am on Sunday. While often situations call for it to be done in other times, the 'huarache' of best times is something like this, from best to worst:

- Early Boxing Day
- Early New Years Day
- Early Christmas Day
- Early Easter Sunday
- Early any Monday public Holiday
- Early any Sunday
- Early any public holiday
- 4am - 6am any day
- 6am - 8am any day
- Before 9am any day
- Morning Sunday or public Holiday
- Morning Saturday
- Morning any day
- Any other time

The good thing about Easter and Christmas/New Year is that a) there are several public holidays close together, and b) it is feasible to plan for extended maintenance windows. All other times it is best to try to limit the window to two hours or less. And you would be surprised how quickly that time evaporates, even with seemingly 'quick' and straightforward work.

So, even though it won't be possible to do some of the physical changes I had hoped, there is still a lot of other work, such as software changes, upgrades, and preparation that can take place. As well as optimisation of the network that always carries some 'risk' and I am loath to do at other times.

Oop.. just got an email from the recalcitrant supplier account manager. For a minute there I thought they may have pulled their finger out. But no. More of 'Your business is important to us, so we have scheduled a meeting next week to discuss the situation', which of course means 'Your money is important to us, and we are hoping to delay fixing any problems until you get fed up or they go away by themselves' - good on ya guys, that's telling it like it is.

Posted by Steve Waddington at 16:28

Wednesday, March 19, 2008

Farewell, Sir Arthur

Sad news to hear that Sir Arthur C Clarke passed away early this morning. I can't remember if it was 'A Fall of Moondust' or 'The Sands of Mars' that was the first books of his I read at around eight or nine. Followed since then, at a rate of two or three a year on average, the other forty odd novels and collections.

Sir Arthur said in his 90th Birthday video that he would like to be remembered as a writer. I don't think any list, no matter how esoteric the opinion of the compiler, would not include him in the 'top three' of SF writers. Sure, you might put Heinlein or Asimov first, or include PJ Farmer, Philip K Dick, or maybe even Kurt Vonnegut. But Dr Clarke is going to be in the three, whichever way you cut it.

To wish to be remembered just for his writing is the mark of true humility of a towering intellect. His contributions to maths and physics may never rival Hawking, Penrose or Einstein. Nor philosophy that of Popper, Hofstadter or Russell. Perhaps some may even argue that his writing was less technically accurate than Egan, less imaginative than Banks and less anthropic than Dick. (Though not me, because those that came after could only follow where he had led).

But how many space craft has any other writer had named after them? Or asteroids? Or species? All geostationary satellites orbit in the 'Clarke Orbit'. Why? Because he invented it, in 1945 working out the principles that led directly to the development of communications satellites.

Space craft where named in his honor during the early 70's Apollo missions. And again in 2001 Mars missions. His inspiration to two generations of NASA engineers is clear. And also to a third generation if the candidate name of 'Clarke Base' for the proposed 2025 Lunar base is any indication.

I don't know how 'inspiring generations' can be measured, but I think that is his greatest legacy to the world

Posted by Steve Waddington at 16:36

Saturday, March 15, 2008

simple, efficient, modular, but scratch simple

A little over four years ago I booted up Exetel's first two routers and saw the first users connect (I think it was eight), accessing the Internet via Exetel.

Two routers was an overkill at the time, the same job could be easily done on one, but of course two gave redundancy if one failed. However the main reason for two was to achieve the separation of functions, which was, and remains, the cornerstone of the Exetel network architecture.

With two routers, one can be the border router and one can be the LNS. The idea was to keep things simple, and if there was a problem, each router 'doing it's own thing' meant the cause would be easy to isolate and track down. It also gave us a modular design that has let us just 'plug in' more routers of the same class whenever extra capacity has been needed, allowing, at least in theory, unlimited expansion of the network within the same operating model.

I say in theory, because of course even the best theory can't take into account the many little adorations that creep into a production network over time.

Since the first two routers were commissioned all that time ago, the network has grown in, roughly, these steps:

- * a second LNS router
- * a third LNS router
- * a dedicated SHDSL router
- * a core router for traffic management and redirection
- * a second border router
- * a fourth LNS router
- * an edge router for ADSL1 and ADSL2 BGP session termination
- * two more LNS routers
- * a second core router
- * a Melbourne LNS router
- * a Brisbane LNS router
- * an edge and two LNS routers for the second NSW POP
- * closely followed by WAN core router for the second NSW POP
- * and a few other smaller routers for special purposes such as firewall or special projects

On average a new router has been added every two to three months, accommodated within the original design criteria of simple, efficient, modular.

So this morning, over my breakfast cup of tea, I was looking into why some routers were experiencing cpu spikes. With six console screens open to various routers and the copious side notes I was making to track what was going on, it occurred to me that 'simple' is no longer integral to the network fabric.

It remains true that any part of the network taken in isolation is fairly simple, and even the route path to any specific destination on the network is also pretty simple and easy to understand. But all those twenty odd routers, their IGP's and EGP's - oh well, all good things must come to an end at some time.

So it's not before time that we will be making the first of the big hardware change-outs to our network over Easter. Starting with the replacement of two LNS's with a Cisco 10000 series router. Once we have that bedded down, the new router will pick up the work of another two LNS routers and an edge router, replacing five 7301 class routers in all. That in itself will much simplify the routing fabric of the Sydney 'A' POP.

Following that will come a second Cisco 10000, replacing three routers at the Sydney 'B' POP. Then will be the phasing out of the core and border routers for replacement with a single, 400Gbps backplane capacity, unit. Effectively, two units will again carry all, or the majority of, the Sydney 'A' POP traffic - but with about 1,000 times the capacity of the first ones we put in. Funny how these things come full circle.

Ok, maybe 400Gbps capacity is an overkill. Sort of like 2Gbps capacity was four years ago when the first 50Mbps of

traffic was switched on. Then again...

Posted by Steve Waddington at 18:30

Wednesday, March 5, 2008

MIPI

I was passed a letter from Allens Arthur Robinson, who, apparently, act on behalf of Music Industry Piracy Investigations Pty Ltd, which, is explained, represents the 'record companies' (doesn't say which ones, so all of them I guess) through their association with ARIA.

It is a followup to a letter sent last March (don't recall seeing that one), requesting cooperation for support for a 'notice and disconnection' procedure. Well, it is sort of a request, along with what I read to be some implied threat, saying that they believe Exetel has a clear responsibility under the Copyright Act to take steps to address illegal file shares on its (our) network. But that we have not yet indicated our agreement to implement the 'notice and disconnection' policy.

Haven't we? How remiss that we have not complied with this legislated requirement. Well, we would be remiss, if in fact it were a legislated requirement. Reading the letter, one might be lead to believe there was some onus to do that, but in fact there is none. Rather, it is a request from one company to another to do something that may very well not be legal.

It seems to me if we have a contract to supply to a customer, someone else saying they think that customer has done something wrong does not provide us with a legal reason to terminate supply to that customer. I don't know, I am no lawyer. Them saying that 'we have evidence that shows that your customer has done something wrong' doesn't really help us either. Exetel being neither a law enforcement agency nor a court of law has very little (legal) ability to interpret 'evidence' (in fact none at all) or make some arbitrary judgment on it.

The proposed policy is 'three strikes and you're out'. MIPI included drafts of two warning notices, followed by a third disconnection notice to be sent by the ISP. The letter goes on to explain that ARIA and MIPI have been trying for some years to agree a code of practice with the IIA, but that those negotiations ended prior to a code being agreed upon. Not surprising really. Without a trial case to set the precedent, I struggle to see how any code could be adopted where the legal ramifications are unclear.

The way I see it is this:

1. ARIA or MIPI gain evidence that shows that a private citizen may be breaching copyright
2. The tell the appropriate law enforcement agency, who investigate
3. The ISP receives a subpoena or other legal request to identify the citizen
4. If there is sufficient evidence, the law enforcement agency brings the case to trial

The way MIPI want to do it is this:

1. ARIA or MIPI gain evidence that shows that a private citizen may be breaching copyright
2. They send the information to the ISP, who then send the first warning notice to the registered user of that IP address
3. ARIA or MIPI gain evidence that shows the same citizen (or the same IP address) may have breached copyright again
4. The information is sent to the ISP, who sends the second warning notice
5. ARIA or MIPI gain evidence that shows the same citizen (or the same IP address) may have breached copyright a third time
6. The information is sent to the ISP, who issues a disconnection notice and terminates the customers service

Think of it like this; someone uses a bus to get to a shopping centre each day, steals some CD's from a record store, then gets on the bus and goes home. The record store sends a letter to the bus company saying that a person is using the bus to get to and from their store to steal CD's. They ask the bus company to caution the person, and if that person

keeps doing it, refuse to take them on the bus any more. Hmm, don't you think the record store owner should just call the police?

Can the bus company refuse service to a paying customer on the say so of a third party? Can an ISP terminate the service of a paying customer on the say so of a private third party? The questions I can think of that would need to be answered are:

1. If the customer is still in contract, can the ISP legally break the contract on the say so of MIPI alone?
2. The ISP has a back-to-back contract with the Carrier. Will the carrier excuse the remainder of the contract and does ending the contract between the ISP and the customer also end the contract with the carrier?
3. Is the customer responsible to pay the early termination of contract fee? Does the ISP still have to pay that to the carrier?
4. If the customer takes action for breach of contract by the ISP, is MIPI liable? Will they pay any damages and costs?

Shaky ground if you ask me. I can't see any proposed code gaining wide support without the legal position of all parties being much more clear than it is.

That in a nutshell is what I understand the proposed code of practice to mean. I don't see how an ISP could agree to support it in practice, not without some very good, expensive, lawyers, and some weight of legal opinion to back it up.

Of course, as far as copyright theft goes, Exetel has taken the hardest stance of any ISP that I know of, and we have done that since we began over four years ago. Not because of any implied threat or stretched interpretation of the Copyright Act. And certainly not because of any direct financial benefit (quite the opposite, from the 'outraged' emails sent from soon to be ex customers). Rather because, actually, it is the right thing to do. We have always made it very clear that thieves are not welcome.

After some trial and error to find the best system, we developed the following process which has worked very well for the last two years:

1. We receive a notice of copyright infringement from a recognised industry source, or their legal representatives
2. We forward the notice to the account holder of the IP addresses identified
3. At the same time, we change the end users public WAN IP address to a private IP address. Any web page they try and browse is redirected to a captive portal that explains why their service has been blocked. In the mean time, and importantly, they are still able to send and receive email, should they wish to use that medium to resolve their issue.
4. The captive portal instructs the end user to resolve the issue with the issuer of the notice
5. Via the captive portal, the end user can select a number of options to resolve the issue, including; terminating their service; doing nothing (in which case the block stays in place); telling us they have resolved the problem with the notice issuer (in which case the block is removed).

I ran this system manually for about eight months before I handed it over to our sysadmins to be fully automated. Over that time it was hard to say just how effective it was, but I did notice the trend that a) repeat IP addresses decreased each month and b) the ratio of new accounts to existing accounts that notices were sent to increased. So it does seem to have a positive effect.

Also, the information is all logged, so should a law enforcement agency investigate, the history of what the end user has told us is there.

The advantages of this system over the MIPI proposal is:

1. It can be fully automated, relatively easily (in fairness, I don't know if the MIPI proposed code can or not, but I know ours can)
2. The end user is directed to resolve the issue with the issuer of the notice

3. The ISP does not breach the contract. It is an option for the end user to terminate the contract if they wish.
4. There is no 'I didn't get the email/notice' or 'I didn't know about it'. If the end user uses the Internet, they will be 100% aware of the issue, guaranteed.

In the event that the MIPI proposed code is adopted by the industry, it should be easy enough to implement. But I do think we are already well ahead of that curve in that respect.

Posted by Steve Waddington at 13:39

Monday, March 3, 2008

Problem Diagnosis with Ping

The two most used tools, and almost always the first used to diagnose a network problem are traceroute and ping. The results they return are however most often misunderstood or interpreted in a way that leads to an incorrect conclusion.

Let's take the ping utility specifically. The common mistake that is made is that whatever the ping result is, is due to the target of the ping. For example, if there is no ping response; conclude that the site is down. Or if there is packet loss or long return times, conclude that it is because of some problem with the target address. While both those outcomes could be the case, far more often than not, they are completely the wrong conclusions to draw.

The common causes of this misinterpretation are:

1. Ping sends a packet to the destination address that typically will traverse several other network points to get there. A problem at any one of those points will cause a non response to the ping query
2. In many cases web sites and other servers sit behind firewalls, and many, if not most, firewalls block ping packets. So while web traffic may reach the site, ping packets may not.
3. The ping packet has a source (the system initiating the ping) as well as a destination, it may be that the source does not have a correct route path to the destination, or that the destination does not have a correct return route path to the source. This could be because of specific firewall rules, an error in the route tables 'somewhere' along the data path, or a specific routing policy deliberately put in place to block access.

The traceroute command can be used to help detect if 1. or 3. are the cause of the problem, which has its own issues, but more on that later. A positive result from either telnet and tcptraceroute will conclusively rule out 2. as a possible case.

Telnet can be used to open a connection any any port, not just the telnet default port. A successful telnet connection where ping has failed is proof positive that a firewall is preventing access to ping packets. Here is an example:

```
$ ping www.cisco.com
PING www.cisco.com (198.133.219.25) 56(84) bytes of data.

--- www.cisco.com ping statistics ---
6 packets transmitted, 0 received, 100% packet loss, time 5008ms

$ telnet www.cisco.com 80
Trying 198.133.219.25...
Connected to www.cisco.com.
Escape character
is '^'.
```

You can see that the ping packet failed, but that telnet to port 80 succeeded in connecting to the server.

So too with tcptraceroute on port 80:

```
$ tcptraceroute www.cisco.com 80

traceroute to www.cisco.com (198.133.219.25), 30 hops max, 40 byte packets
 1 192.168.6.254 (192.168.6.254) 8.557 ms 10.624 ms *
...
15 www.cisco.com (198.133.219.25) 289.162 ms 237.972 ms 242.171 ms
```

Another common error using ping is that the results of just a few ping tests are indicative of the condition of a data path. It may be true, but such a conclusion can only be relied upon over a statistically meaningful sample size. Also, to be

truly accurate, the distribution of packets responses outside the acceptable level needs to be known.

For example, as single ping test of four packets where one packet is dropped, can not, in any meaningful way, be used to conclude that there is 25% packet loss on that circuit. Ten thousand ping tests, over several hours where there is say 5% lost has far more meaning; however consider if the test was done over 24 hours, and for one hour the target site was down. The 100% loss during that hour looks like a general 5% packet loss over 24 hours.

It is therefore important to review the record of the ping test and see if the distribution of any packet loss is regular or confined to a specific period, before a real conclusion can be drawn.

A third common error is that the cause whatever is result is gained is due to the target site. For example, say 5% packet loss was found when pinging www.3com.com, this by no means indicates that the problem lays with that site, rather, the problem could be with any of the points along the data path to that site, inclusive the source (my own computer):

```
$ traceroute www.3com.com
traceroute to www.3com.com (192.136.34.41), 30 hops max, 40 byte packets
 1 192.168.6.254 (192.168.6.254) 10.285 ms 13.316 ms 14.440 ms
 2 129.1.233.220.exetel.com.au (220.233.1.129) 132.994 ms 135.387 ms 136.312 ms
 3 241.0.233.220.exetel.com.au (220.233.0.241) 137.192 ms 141.296 ms 162.018 ms
 4 10.0.1.1 (10.0.1.1) 168.530 ms 174.358 ms 176.908 ms
 5 38.2.233.220.exetel.com.au (220.233.2.38) 177.729 ms 188.233 ms 189.122 ms
 6 359-ge-0-0-0.GW5.SYD2.ALTER.NET (203.166.92.57) 197.691 ms 85.598 ms 156.625 ms
 7 0.so-0-2-0.XR3.SYD2.ALTER.NET (210.80.33.189) 158.108 ms 159.430 ms 160.260 ms
 8 0.so-4-3-0.IR1.LAX12.ALTER.NET (210.80.50.249) 305.124 ms 305.952 ms 306.775 ms
 9 0.so-5-0-0.IL1.LAX9.ALTER.NET (152.63.48.65) 313.518 ms 321.047 ms 321.868 ms
10 0.so-5-0-0.XT1.SAC1.ALTER.NET (152.63.0.98) 405.111 ms 406.359 ms 407.241 ms
11 GigabitEthernet6-0-0.GW9.SAC1.ALTER.NET (152.63.55.73) 331.091 ms 337.600 ms 341.527 ms
12 eds-gw.customer.alter.net (63.114.61.154) 357.930 ms 287.765 ms 310.755 ms
13 205.141.209.3 (205.141.209.3) 311.606 ms 312.502 ms 313.587 ms
14 10.231.1.2 (10.231.1.2) 341.277 ms 342.101 ms 342.931 ms
15 205.141.209.133 (205.141.209.133) 344.380 ms 345.861 ms 346.689 ms
16 ip-192-136-34-41.ip.3com.com (192.136.34.41) 261.317 ms 266.998 ms 346.689 ms
```

You can clearly see the number of hops the data must traverse. In this case there is no evidence of any problem along the data path. But if the traceroute looked like this:

```
$ traceroute www.3com.com
traceroute to www.3com.com (192.136.34.41), 30 hops max, 40 byte packets
 1 192.168.6.254 (192.168.6.254) 10.285 ms 13.316 ms 14.440 ms
 2 129.1.233.220.exetel.com.au (220.233.1.129) 132.994 ms 135.387 ms 136.312 ms
 3 241.0.233.220.exetel.com.au (220.233.0.241) 137.192 ms 141.296 ms 162.018 ms
 4 10.0.1.1 (10.0.1.1) 168.530 ms 174.358 ms 176.908 ms
 5 38.2.233.220.exetel.com.au (220.233.2.38) 177.729 ms 188.233 ms 189.122 ms
 6 359-ge-0-0-0.GW5.SYD2.ALTER.NET (203.166.92.57) 197.691 ms 85.598 ms 156.625 ms
 7 0.so-0-2-0.XR3.SYD2.ALTER.NET (210.80.33.189) 758.108 ms 759.430 ms *
 8 0.so-4-3-0.IR1.LAX12.ALTER.NET (210.80.50.249) ** 806.775 ms
 9 0.so-5-0-0.IL1.LAX9.ALTER.NET (152.63.48.65) 813.518 ms * 721.868 ms
10 0.so-5-0-0.XT1.SAC1.ALTER.NET (152.63.0.98) * 1406.359 ms 1007.241 ms
11 GigabitEthernet6-0-0.GW9.SAC1.ALTER.NET (152.63.55.73) 731.091 ms 737.600 ms 1341.527 ms
12 eds-gw.customer.alter.net (63.114.61.154) 357.930 ms **
13 205.141.209.3 (205.141.209.3) 811.606 ms 812.502 ms 813.587 ms
14 10.231.1.2 (10.231.1.2) 741.277 ms 742.101 ms 1342.931 ms
15 205.141.209.133 (205.141.209.133) * * 746.689 ms
16 ip-192-136-34-41.ip.3com.com (192.136.34.41) 761.317 ms 866.998 ms *
```

It would be reasonable to conclude that there was some serious problem between hop 6 and hop 7 that is causing the ping test to return its lossy result.

To conclude, we can see that ping:

1. is a useful tool to indicate where a problem may be

2. should be used in combination with other tests to eliminate false positives
3. should not be used for small, isolated tests
4. is a good indicator of problems over statistically meaningful sample sizes

Posted by Steve Waddington at 12:56

Tuesday, February 26, 2008

When is a Profit not a Profit?

I usually stick to commenting on things I know about, or at least like to think I know about. But I was sent one of the larger public company ISP's half year report, which I find interesting as general indicator of the 'health' of the non-Telstra ISP world in Australia.

Wow, \$9M profit for the half year and \$25M EBITDA, that looks good. Even a small interim dividend. Things sure look rosy.

But wait a minute, cash is down \$2.3M. Doesn't a profit mean cash should increase? You can tell that as a network engineer, I make a pretty poor accountant.

Let's see, net cash flow in: \$32.4M, Net cash flow out: \$9M. Therefore, profit = ~ \$23M. I get that. So how can cash go down? Ah, repayments on borrowings: \$19M and dividends: \$6M, leave a \$2M cash shortfall. It seems to me like someone sitting down at a poker table with \$1,000, expecting to lose it all, and therefore if they leave with \$500, then they call that \$500 'profit'.

As I say, I am no accountant, so what do I know of the technicalities of what can be reported as a profit and what can't. I don't know if such complex company reporting can be equated to a household budget, which I imagine in the that ISP's 'home' would run something like this:

Husband: Honey, I got paid \$10,000 this month and after clothes for the kids, school fees, petrol, groceries and other bills, we have \$6,000 left for us.

Wife: That's wonderful sweetie, you are such a great provider and so clever

Husband: (beaming with pride) gee thanks honey, you are pretty clever too, picking a man like me. Let's go out for dinner tonight, we deserve it.

Two days later

Wife: Um, sweetie, the bank called today and said our cheque account was overdrawn

Husband: I don't see how that can be, even after our dinner the other night, we still had \$5,500 in the account. I'll call them now and sort it out

Husband (on phone): Yes, uh huh, I see. Ok, thank you.

Wife: Well?

Husband: Ah, well, you see honey, we had \$5,500, but the car payments and the mortgage add up to \$5,900, and then dinner was \$500, so we are about \$400 overdrawn.

Wife: You idiot. Why didn't you take our interest payments into account?

Husband: But I was just working out our profit before that. You mean I have to count all our outgoings?

Wife: How did I ever marry such a moron.

Husband: Now honey, don't be like that. It's not like this is the first time. Remember last month and the month before that? In fact every month we have been married. Can't you just dip into your personal savings like you usually do to make up the difference? Honey? Snookums? Sweetie?

Wife: Well.... ok, just this one time more....

A month later

Husband: Honey, I got paid \$10,000 this month and after clothes for the kids, school fees, petrol, groceries and other bills, we have \$6,000 left for us.

Wife: That's wonderful sweetie, you are such a great provider and so clever

Husband: (beaming with pride) gee thanks honey, you are pretty clever too, picking a man like me. Let's go out for dinner tonight, we deserve it.

....

Posted by Steve Waddington at 12:11

Friday, February 22, 2008

Whitepaper site up and running

The whitepaper wiki is now operational and we are starting to populate it with some of the documents written for Exetel in the last four years.

<http://whitepapers.exetel.com.au>

I am not sure MediaWiki is the best platform for this. Its security features are fairly basic and clunky. I don't know if it is common to all Wiki platforms, but the editing features are very primitive.

But, it is a start. Changing the platform will be trivial if a better one is found.

Posted by Steve Waddington at 12:40

Thursday, February 21, 2008

P2P Cache work and core network plans

We added a third border router today which should, in theory, remove any doubt that the lower than expected performance was caused by the capacity of the existing two routers.

We are also looking at future options for our core and border routers. One suggestion made by PeerApp was that a Cisco 6506 with a 720 route engine would be a good option to replace our, now, three 7301 routers. While they are excellent switches, I have in the past found the route engines to run out of steam much more quickly than expected. However, the last time I used one was in 2003, so I imagine the performance profile may have been improved since then.

Certainly the pricing we have so far is very attractive - less than \$30k for what would essentially be a 16 GE port router.

Another possibility is the Foundry XMR switch/routers. The specs on those are very impressive. It will be interesting to see how the price compares.

To round out the options, we are also considering Juniper in the mix. My feeling is that price/performance will be good, but the performance will be vastly more than we need, and therefore the cost will be more than we can justify. We will have to see.

I have been thinking about making this change for the last eighteen months, but am very reluctant to move away from a system that has proved so incredibly robust for such a long time. Zero downtime in four years is sort of hard to beat.

But maybe now the time is nigh. Meshing BGP on the border and core routers is a handful, and we will have to move to a route reflector on the next addition.

Posted by Steve Waddington at 13:22

Wednesday, February 20, 2008

Travel tips for the US

The US in general

Tip. Most people in service industries work for the basic wage of \$6.25 per hour, and rely on tips to make a living wage. If you can afford to travel to the US, you can afford to tip. 15% is 'normal'

Accent. Most Americans can't pick an everyday Australian accent, they expect all Australians to have a broad accent like Paul Hogan or Steve Irwin. American's are pretty xenophobic and don't like non-American accents, so if you see them lose interest or 'shut down' when you speak to them, try adding in a 'no worries mate' or some other recognizable Aussie phrase. If they know you are an Aus-see, they will be much more open.

Food. An 'entre' in the US is the word they use for the main meal, and they call the entre a 'starter'. The servings are huge. A starter is as big as a really big main meal, and a main is about twice that size. Diners tend to offer the best value of nice food for a reasonable price. You can get breakfast for about \$10, which will be more pancakes, butter and syrup than three people could reasonably eat, or a huge mound of toast, bacon, eggs, hash browns and grits, and all the coffee and juice you can drink - you won't need to eat again until dinner. Never, ever eat at a Denny's restaurant, it is the worst possible food. Many of the other food chain stores are quite reasonable however.

Domestic Flights. All flights are over booked. Just because you have a ticket is no guarantee you will get on that flight. Either get to the airport three hours early, or do your own seat assignment on-line the night before you travel.

Watch out you are not carrying 'weapons of terror' with you, which apparently include: plastic paper scissors, a school ruler, a calculator with flat batteries, a pencil sharpener (has a small, though sharp, blade), or anything that has, or could be broken to have, a sharp(ish) point or blade - they will all be confiscated. Strangely though, pens and pencils are ok.

Getting around. For one or two people the cheapest way to get from any airport to your hotel is via a Super Shuttle, which operate out of all airports and cost about \$12. A Taxi will be about \$55 in most cases. Also very good value is car hire, which is about \$35 per day for a Commodore/Falcon equivalent, or \$40 per day for a Mustang convertible, or \$50 per day for a Cadillac (6.8l V8) or Lincoln Town Car (7.0l V8).

There is in theory a 55mph (88kph) blanket speed limit in the US. Forget it. On the freeways the traffic goes at its own speed, and if you are not going at the same speed as the traffic, which is anywhere from 70 to 90mph (120 - 150kph), then you are a danger to yourself and other vehicles, as cars and 18 wheeler trucks swerve out from behind you at 80 miles per hour because you are in their way.

Your Australian drivers license is good for car rental in the US. Alamo or National usually have the best deals in most cases.

Duty Free. Forget about buying duty free liquor at the airport, or anywhere. Just go to a Wal-Mart or Wal Greens, it is about half the price you would pay 'duty free'.

In Las Vegas

Hotels. If possible, try and time your trip so you are there mid week. Weekend rates are \$350 for a nice room, mid week rates are \$50 for the same room. Set the Luxor as the minimum quality hotel, if you can get a good rate at Mandalay Bay, MGM Grand, the Wynn, Caesars Palace or the Bellagio, then any of those will be fine. You will have a much nicer time if you stay at the better hotels.

Food. Every hotel has one or more buffets costing from \$5 - \$15 and usually having two or more 10 metre benches with amazing food. Check google for 'Las Vegas best value', there are plenty of sites that tell you the best buffets and restaurants.

Many hotels also have specialty dishes they are famous for. You MUST try the ribs at the Stardust or the Hilton, or a \$5 steak at Binions, or the pasta at the Venetian, or the shrimp buffet at Margaritaville, or Chinese at (of all places) Paris, a Caesar Salad at Caesars (of course). The Luxor is mean to have very good steak, though pricey (\$50), though I

haven't tried it.

Drinks. If you should indulge in any of the games on offer, then a pleasant cocktail waitress comes around every half hour and will get you any beverage you want, for free. Of course, you should tip her a buck for each drink.

If you are going to pay for a drink, then plenty of places offer yard-glass margaritas for \$1-\$2, or other enormous amounts of drinks for not much. Then try finding your way out of the Casino.

Soft drinks, tea or coffee at most places are 'bottomless' - you pay once and if you keep drinking it, the waiter will keep filling it up.

Games/Gambling. That is what after all pays for all those hotels and attractions. As soon as you step out of the arrival gate at McCarren airport, there are poker machines.

All the hotels are designed so that to get from the reception desk to your room, or out of the hotel, or across the street, you have to walk through the casino floor.

In Las Vegas, it is also the 'done thing' to tip the dealer, \$1 for a small win or \$5 for a big win seems about the right amount.

For poker, the rake on small stake games is much better than at Australian casino's and the games are better structured. \$1/2 games are soft and offer good value.

For blackjack, most casino's on the strip now only give 5:6 for blackjack, which makes those games expensive to play. Off strip plenty of places still pay 2:3. However there are \$1 and \$5 tables, and the drinks are just as free as anywhere else.

For roulette, there are 50c and \$1 tables, as well as \$5 and up. However all US casino's have '0' and '00', so the house odds are 2 in 36 rather than 1 in 36 everywhere else in the world.

Unless you go mad, \$200 should last you 10-20 hours at low stakes games, get you all the drinks you want and most likely a comp voucher for a free meal or two at the buffet.

Shows. All the best acts in the world play at Las Vegas pretty much all the time. If you have to pick one, then see a magic show. If you can pick a couple then just look for the best rated ones via a google search.

Other things to see/do. What, apart from all that? Here are a few:

- The fountains at the Belagio (free, watch from the street)
- The pirate show at Treasure Island (free, watch from the street)
- The Star Trek rides at the Hilton - totally amazing, better than Disneyland
- Quarks bar at the Hilton
- The rollercoaster at New York New York
- The StratosFear rollercoaster at the Stratosphere - it winds around the top of a 30 story building
- The M&M arcade - 4 levels of M&M merchandise, but M&M's are only 25c a packet, and it has a reasonable 4D ride.
- The light show at Downtown Las Vegas (free)
- The canals and the wax works at the Venetian
- Just walk down the strip, a sight in itself

Shopping. There is a super mall about 2 miles out of Las Vegas. Clothes and shoes will cost about half what you pay in Australia.

Any Borders or Barnes and Noble store will have hardcover books at the price of paperbacks here - about \$18. Paperbacks are \$10.

Any record store or Virgin Megastore will sell CD's and DVD's at half the price you pay here. Just make sure your DVD player can play region 0 or region 1 DVD's.

You can pick up any software title for half the price you would pay here.

The BEST place to go is Fry's, near the supermall, just out of town. Imagine a shop the size of the biggest Bunnings

warehouse, but instead of hardware and home goods, has isle after isle of hard drives, mother boards, peripherals, software, monitors and any computer or network component you can imagine. Just load up your shopping trolley (literally) and hope your credit card will make it.

Near by. 1 hour drive from Las Vegas is the Hoover Dam. Take the tour, it is well worth it.

2 hours drive is Death Valley. Not much to see, but worth the detour of you are going that way anyway.

There are also half day tours to Red Rock canyon, which is meant to be impressive. And full day tours to the Grand Canyon. I have only flown over the Grand Canyon and even at 30,000 feet it is very impressive.

Posted by Steve Waddington at 16:11

Friday, February 8, 2008

Carrier License

Our application for a carrier license was approved today.

A few small changes were needed on the application. But Just goes to show what can really be achieved when a competent, efficient, busy person is put on the job.

Well done Larry. I think that must be some sort of record for least time spent to obtain a result. Today, there is joy in Mudville.

Posted by Steve Waddington at 11:55

Security Mostly A Waste

Ha, knew it:

http://www.darkreading.com/document.asp?doc_id=145224&WT.svl=news1_1

For a commercial production network, security should only be viewed as a means to make the network run more efficiently and save time. Inefficiency and time wasting being factors of having the network compromised and the unnecessary time needed to fix it back up.

Every company has a duty of care to protect customer and proprietary information - but in a cost effective and efficient way. Once you have a security 'department', you can be pretty sure they are going to want to justify their existence in much the same way that other superfluous dinosaur of a previous century, the marketing department, does. By spending lots of money and needing ever more resources, in return for some very nebulous, and never directly measurable, results.

The best security I have found comes from a competent, busy, engineer, who understands that a breach of the network will cost him more time than a sensible and efficient security policy. Those people are able to very effectively gauge the fine line of security that works and is not cumbersome, opposed to security that works but is restrictive and/or over complicates operations.

It seems to me that many corporate IT dollars spent on security come from FUD (which few security departments I have seen are above helping along). The unfortunate thing is, there are many not-so-competent engineers and admins, who accept 'out of the box' security as enough, and thus provide the fodder for the horror stories we hear about from time to time.

Posted by Steve Waddington at 11:29

Tuesday, February 5, 2008

Whitepaper Wiki

The whitepaper site is at 'alpha' stage at the moment. Having selected wiki as the publishing platform, there now seems to be an issue with controlling how editing is done while still making it available for wider public viewing.

Something I am sure the brains trust of our sysadmins will overcome.

John has published several dozen papers on management - very interesting reading, and a 'refresher' for me on some of the more subtle aspects of our business process.

In the meantime I am preparing the first five or six papers, mainly focusing on network management and BGP routing, which I am taking from technical briefing notes I have given in the past.

Another half dozen or so will come from our noc-list archives. Looking back to the first post in September 2004, I had forgotten just what a rich source of information, and record of progress, that is. It is easy to see how so much of the international collaboration that has built the global Internet has been facilitated with email lists. Email lists would have to be the most under-rated killer-app there is.

Posted by Steve Waddington at 13:36

Monday, January 28. 2008

Taming of the Spike

We moved the NetEnforcer to the inside interfaces of the UB2000 cache because we decided it would be the best way to control the midnight spike of traffic. It also let us ensure that there was no policy error on the NetEnforcer that was stopping the cache from getting its full 'suck' from the Internet.

The result was successful, as you can see on the weekly graph:

Though it is more obvious on the monthly graph:

The extra 4Tbyte storage was installed eight days ago and has been slowly filling up:

Volume	Size GB	Used GB	Avail GB	Use %	State
1	397	360	37	90.83	active
2	397	366	31	92.38	active
3	397	361	36	91.00	active
4	397	358	39	90.24	active
5	377	343	34	91.06	active
6	377	353	24	93.74	active
7	397	158	239	39.86	active
8	397	367	30	92.61	active
9	397	359	38	90.62	active
10	396	358	38	90.59	active
11	794	209	585	26.34	active
12	794	272	522	34.32	active
13	757	256	501	33.82	active
14	794	251	543	31.70	active
15	794	236	558	29.74	active
Total	7862	4607	3255	58.60	15/15

And I think our generated bandwidth is improving - although it seems a glacial speed. We are now seeing about 193Mbps average generated bandwidth per day. The PeerApp folks think it will be two weeks before we see the full benefit.

We really need to get to 300Mbps to make the cost/benefit irrefutable. Particularly now we have firm offers of bandwidth for 2008 well under \$200 per mbps.

Posted by Steve Waddington at 16:39

Friday, January 4, 2008

PeerApp Update

We had a conference call with PeerApp this morning focusing on the lower than expected traffic generation from the cache. We are feeding it an average of 620Mbps of cachable traffic (selected from 1,300Mbps), but only seeing 200Mbps average generated traffic. PeerApp had analysed the performance and come up with a couple of things. First, that the disk storage was a bit on the low side at 4Tbytes, and we should have another 4Tbytes which they expect would allow about another 100Mbps of generated traffic. I remember in 1997 installing a 1Tbyte Network Appliance disk array - an unimaginably large amount of disk space in those days. We were particularly pleased with the price we paid (and Network Appliance were particularly displeased) at only \$400,000. Anyway, another 4Tbyte unit is on its way to us. That must be, oh, \$2,000 worth of disks now I guess. The second thing was that in other countries, PeerApp see cache lives of 2-3 weeks for popular files. On our cache however, they are seeing a cache life of 50 hours at most. They thought that may have something to do with the broad multiculturalism in Australia, compared to say, the relative mono cultures of Thailand and Portugal. Though I didn't fully understand why that would be. But I wonder if it doesn't have something to do with the writers strike in the US. I suppose we will see over time. We also talked about the 'big spike' we get at midnight each night. PeerApp think it could be because of one or more bandwidth policy on the NetEnforcer. However we have given the cache IP address block an unrestricted access policy which is the same as if it were wired directly to external bandwidth. I think the reason for the spike is because the 'uncounted' period starts at midnight, and so many people start their downloads at that time. The spike is caused by the most popular cached files being delivered quickly in the first couple of hours. The quadratic decline of the curve suggests that explanation as more likely to me than a bandwidth management policy. PeerApp are kindly going to analyse and bench test our NetEnforcer config with the folks at Allot anyway, so it will be interesting to see what they can come up with. I am sure there is some optimisation that can be done there in any event.

Posted by Steve Waddington at 11:42

Wednesday, January 2, 2008

Back up your email

I have an answer to this 'problem' reported by Computerworld. People who sell hard drives aren't going to like it, but still. My solution is this - just delete it. Do you really need every email and every piece of data from the last fifteen years? 'Ah' I hear all the habitual hoarders say 'but you never know what might be needed one day'. It's true, you never do. But does anyone ever stop to consider exactly what it might be needed for? Let me tell you, the chances are about 999 to 1 that whatever that ancient data is going to be used for, isn't going to be to anyone's benefit in the present day. Consider; "Mr Doe, this email you replied to in 2001 indicated clearly you were aware of the problem, did nothing to correct it, and therefore are fully culpable for the consequences. Can you tell the court why that was?" Assuming Mr Doe could even remember the email, or what it was about, what are the chances of anyone else ever understanding the context, the non-email responses from other people (for which there is, obviously, no record), or any other of the thousands of things that would have continued to the particular comments made in the email? Yet, there is a written record, undeniably correct in every word, as evidenced by the multiple copies of the people it was sent to on the server backup disks. The fact is, as an undeniable trail of blame, email is at best as reliable as the spoken word, and no more. While a verbal contract is, I have heard, just as binding as a written contract, should email even be given the same weight as a verbal agreement? I don't think it should. The informality of email (as the standard email salutation of 'Hi ', shows) causes it to be treated by the vast majority of people, personally and in business use alike, the same way a casual conversation might take place at a bar or at a party where you don't know anyone. Yet unlike those casual spoken remarks, the archived storage of every word you every typed to anyone ensures your liability long, long into the future. And who knows in what context it might be viewed then? Poor Mr Doe. Eleven years for you mate. Contracts, deed titles and other important business documents should certainly be kept. Sure, scan them into some digital format and burn them to some un-erasable media. Save a few trees. No problem. But is email really an 'important business document'? If someone of a litigious mind convinces a court it is, what hope is there of defending what was undeniably said so long in the past, when the context and other circumstances are so far beyond anyone's recall now. So delete it. Keep email for about as long as a verbal conversation could be remembered - a month at most (a day in my case). Anything important should, and will, go into a contract or some other important archive worthy document any way. Anything else is either current or done with. Either way, it just isn't needed. Isn't the fourth circle of hell reserved for hoarders (along with their counterparts, the wasters)? I am sure Dante would have reported some mighty big storage arrays there if he knew what he was looking at.

Posted by Steve Waddington at 17:09

Tuesday, January 1, 2008

Direct Bandwidth in to Melbourne

Melbourne and Brisbane now both have some Internet bandwidth delivered via a directly connected circuit rather than from Sydney. Brisbane has been on for just over a month. Melbourne was switched on on Friday, along with a two of other circuits for HSPA and ADSL2 carrier cross-connect.

After a bit of tweaking over the weekend, we have the Sydney-Melbourne traffic balanced to within 5Mbps of optimum - 5% in other words. It will be interesting to see if any difference is reported or commented on by users there. My expectation is it may make a slight difference to some traffic local to Victoria, Adelaide and Perth (which some carriers hub back to Melbourne). But generally there should be very little difference, since most other national traffic is hubbed from Sydney and the Sydney-Melbourne hop adds very little to other traffic that may come into Melbourne directly.

It does let us reduce the cost of interstate Internet bandwidth by negating the backhaul cost to Sydney for additional bandwidth. Plans are in place so that over time, as we increase the direct bandwidth, autonomy will also increase in each POP.

Next for direct bandwidth will be Perth. Which I would like to see in by the end of September.

Posted by Steve Waddington at 09:31

Sunday, December 23, 2007

Whitepaper Site

John mentioned something in passing the other day that struck a cord with me; which was to publish the internal whitepapers over the last four years as part of the Exetel web site. As an avid seeker and reader of whitepapers on the Internet, I know first hand how useful they can be, in two ways. First, they more often than not contain ideas and methods of real use, or if not that, the germ of an idea that can be put to use. Second, they are, like blogs, a very good way of getting a company point of view across. And the third thing they do, if indeed they do contain information of value to others, is increase the reputation of the author. Which, for the people at Exetel, would be very well deserved in the case of many of the innovative and unique things we have developed. Some of the obvious ones that spring to mind are:- Andrews failed login and blocked access diversion method- Vinna's bandwidth management method- Simone and my virtual training program - The support coefficient method of measurement James used- My traffic splitting and 'zero cpu impact' bandwidth management system - Martin and James' method for remote support management- Raymonds 'beyond carrier grade' bill checking system (perhaps the first accurate carrier billing system ever?), that enabled us to claw back hundreds of thousands of dollars in over charges. And that's just a few that I can quickly recall in a few minutes. Really, at every part of the company I look at I can see something that is better than any other thing anyone has done. So, quite rightly, mellow yellow aside, everyone should have their 5 minutes of fame, and more, for their contributions. John and I spoke further on that topic a couple of days ago, and the true power and meaning of open source. Love, peace and mung beans for some, no doubt. Yet a powerful way of achieving real commercial benefit, when it is understood, for others. You take what you want, you give what you can, everyone benefits. Socialism? Hah! I think not. Open source comes from real people, who can produce real things. They live and die on the usefulness of their ideas and code. You want more of what they got? You pay for it, the amount? Whatever you think it is worth (or not, in which case, don't expect too many more releases). You want support, or to commission more work from the author? No problem, here is the rate sheet. Open source says; Here is it, my work. Let the market decide. It seems to me most like the Dutch Masters of the 17th century, that the creative output is done in the name of self interest, while providing a service other want. Capitalism, as it should be. So too maybe it is time Exetel put our ideas 'out there'? (As a matter of fact, we have been doing that for a long time. John has daily documented, in detail, almost every aspect of the commercial operations of Exetel over the last four years, for anyone who has cared to read it). I'll talk to Andrew and Raymond over the next week or so, and see what we can come up with.

Posted by Steve Waddington at 12:41

Saturday, December 22, 2007

Photons and Electrons

I heard a good one yesterday, and I thought I had heard them all. We have been having the usual dramas of getting a new GE circuit installed that we ordered about 10 weeks ago. Nothing unusual, but here is a list to bore anyone who isn't already totally bored with this sort of thing:

1. The circuit was eventually completed only four weeks after the due date
2. However, when the co-lo engineer went to run a cross-connect to it from our rack, he couldn't find it
3. The telco told us it had definitely been installed
4. We sent the engineer back with clearer information, still couldn't find it
5. More to-ing and fro-ing, Telco tells us they mixed up the job numbers and it was another circuit. Ours will be installed tomorrow
6. Send the co-lo engineer back the next day to complete the cross-connect, still can't find it
7. More to-ing and fro-ing, we were given the wrong rack location by the Telco, now we have the right one
8. Co-lo engineer finds telco hand off is multi-mode, not single-mode as expected. Co-lo has no MM fibre, will have to order in, expect 10 days lead time
9. Co-lo agrees to put in a multi-mode to single-mode converter and completes cross-connect
10. We connect to our router and test, all is good
11. Exetel engineers on site to prepare for migration to new GE circuit, find circuit is dead
12. We replace FO transceiver on our router with no change and open fault with Telco
13. Telco sends tech to exchange to trace circuit problem. Engineer disconnects circuit, tests for an hour, and reports no fault and can see customer end equipment is live
14. Seems odd to us, since our equipment is not yet connected
15. Telco realises they sent tech to wrong exchange and tested wrong circuit (we heard the circuit was for another ISP that was put off the air for an hour)
16. Next day, telco send tech to right exchange and starts testing the correct circuit. Tech traces the circuit to the co-lo MDF and still finds no fault. Believes problem is with the co-lo SM to MM converter

And this is the one I have never heard before: The Telco tech said something like "Converting from Multi Mode to Single Mode often causes a problem because of electrical impedance differences between the fibre optic circuits" Now, as I understand it, and I could be wrong, the conversion of photons to electrons follows $E=mc^2$, and current media converter technology is not yet able to change energy into mass. Maybe he was referring to some other property of the alpha function, as yet unknown to physicists, or anyone other than Telco field techs? Or maybe it was just BS he thought the customer would believe that would expunge him from needing to do any further work on the afternoon of the last working day before Christmas? I guess we will never know. Anyway, now I am pretty sure I _have_ heard them all. Needless to say, we didn't accept that explanation and requested another field tech be sent. We also go the co-lo engineer on to the job to see if their cross-connect could be the problem

To cut this already long story a bit shorter; the Telco tech and co-lo engineer, along with our engineer over the next four hours check every part of the circuit and equipment from the Telco hand off to our CPE. In the end a loopback was connected to the fibre termination in our rack, and suddenly, the circuit interface lights went green. The loopback was disconnected and the circuit connected to our CPE, which also showed an up/up connection. Like magic it was all now working, yet 'no fault was found'. Amazing, but true. I have been running large and small packet tests over the circuit for the last 12 hours - over 3 million packets now without an error. So I think the circuit is good. I am really glad the guy in telco circuit provisioning remembered to flick the 'on' switch for the circuit in the end. But really, I shouldn't be so cynical this close to Christmas.

Posted by Steve Waddington at 14:11

Friday, December 21. 2007

2nd Alteon Redirector Installed

The eagerly awaited 2nd Alteon switch cleared customs and was installed by Vinna last night. PeerApp worked on the config over night to smoothly integrate it into the system. One more thing to do though - we need a second GE cross-connect run from the border router rack to the server rack. It looked like we had a spare since we could locate a free 'end' in each rack, but there is no circuit there on testing. But, we should have the cabling sorted today in any event. And then... the testing!

Posted by Steve Waddington at 09:17

Saturday, December 15, 2007

The Smart/Stupid Syndrome, and Occam's Razor

It is something I have noticed over the last 20 years. Perhaps 'stupid' is too strong a word, but by stupid, I mean people who don't really think things through, do them by rote, or perhaps just do what they are told to do rather than what should be done, because they have been told they are not allowed to change things. So by stupid I mean people who, for one reason or another, perform stupid actions, rather than are actually stupid themselves. Here is what happens; the stupid ones use complicated code, or expect and look for complicated answers, because they don't know any better, and I guess things do in fact look complicated from that perspective. But then smart ones, not to be outdone, add even more complexity and look for even more complicated answers and solutions to problems. Their rationale seems to be that for problem 'x' if 'stupid' only sees x (complexity), then the answer must be at least xy (complexity) where $y > 1$. Or perhaps they are just trying to move the answer to a more rarefied level of obscurity until they are the only ones (being the smartest people around) who can understand it, and therefore will not be bothered by the stupid asking them stupid questions. The early warning signs, of what turns into a career long trend for many, is if you ask the question 'when will it be done/ready/available?' instead of a date and time, receive a torrent of information about all the reasons it can't be done, won't be ready, and/or the many problems encountered or likely to be encountered along the way. In a junior engineer it is bad - it leads to frustration for the customers they deal with and unnecessary delays getting it fixed. Any sign of 'smartness' needs to be stamped out and replaced with solid engineering practicality as early as possible. In a senior or engineering manager, or even CTO, it is a total disaster for their company. I have seen it several times now, where the smart prick (excuse me, but that is the word that has to be used) drives their company into the ground with huge layers of complex 'solutions' that no one can understand (except them because they are smart, right), is as saleable as the Emperor's new clothes, and impossible to manage. It is too late for them, and all too often their company and its investors. On that basis, I must be, and I am very happy to be, stupider than stupid. Because, I don't care what anyone thinks; for problem 'x', remove complexity and look for the solution of exactly x or less - ie xy where $y < 1$. And if I can't do that, then I consider the very real possibility I am on the wrong track. Funny thing is, it has always worked. Without exception. Sometimes I slip (less now than I used to), and think I am smart, and have ended up with quite large and unwieldy solutions. Which, although they mostly worked, have always been wrong, and have always ended up being replaced by something much simpler and far better. The other funny thing is, I have just about never had a problem 'selling' the complex solution, but I have always had a very tough time selling the simpler one. I think most people, engineers in particular, like complicated and involved things. How else can we get the accolades and awe our technical prowess rightfully deserves? However, 'entia non sunt multiplicanda praeter necessitatem', the law of Parsimony always applies. I have to say one of the best things about the last four years is working with partners that wouldn't accept complex solutions in any event. Occam's razor is wielded ruthlessly, as it should be - just a bit of engineers blood on the boardroom walls from time to time, nothing to see here, move along. And so one more pitfall of a growing business is avoided by Exetel.

Posted by Steve Waddington at 09:51

Friday, December 14, 2007

Cache Drama

We completed the physical change needed for second supplier traffic to be routed via the cache on Wednesday. Then requested PeerApp to make the changes needed redirect a portion of the total traffic symmetrically to the cache. We only wanted to send about 1/3 or a little more to the redirector switch, until the second switch arrives, because of the already high CPU load. The changes were made on Thursday at 5:30pm, ready for the evening load increase. The results weren't exactly what expected, but, the indicators at the time showed traffic and loads within acceptable parameters. All seemed well. Until...Midnight, and the 'free' uncounted download period started. That is usually when we see a big spike in traffic anyway, but this time there was an insidious error in the ACL code. You can see the problem here: `access-list 2009 permit ip any 220.233.128.0 0.127.255.255` Obvious, isn't it. So instead of sending 1/3 of our traffic to the redirector, 2/3 was being sent. Which was ok - because it wasn't much different to the volume being sent that way before - no extra stress and no flag-able error conditions at that stage. Unfortunately, the same error was present on the access-list responsible for load balancing traffic inbound to our routers across 2 1Gbps circuits. But instead what happened was all traffic was directed down a single circuit. Fine until midnight when the net traffic was 920Mbps. Not so fine after midnight when traffic goes to 1.3Gbps. The problem in diagnosing the sudden contention issues was three-fold: 1. Higher traffic was expected anyway at that time, and individually, nothing was too far out of normal parameters to clearly point to the cause 2. The plethora of new measurement points made the source of the problem difficult to pinpoint 3. Ping and other low port traffic was pretty much unaffected, so the 'standard' troubleshooting tools showed nothing out of the ordinary in the test cases run to figure it out. Every network manager likes a bit of excitement in their life - not. The dramatic irony, as always, is that after hours spent trying to figure it out, and an emergency workaround to give temporary relief of the problem, the solution was a change to just two lines of router code. One minute to correct and apply. Ok, drama over and it is all good this morning. These are the results we are seeing sending, currently, 307Mbps of symmetric traffic to the cache: Which is 42Mbps of generated traffic for the 300Mbps offered. Less than I was expecting, so we will need to look into the reasons for that. However we are seeing a big improvement to redirector CPU (note, the timestamps are out on the graphs, I think the timezone on those servers is for Hong Kong): No sign of the 2nd Alteon switch yet, so it will be well into next week before we can get results for all traffic.

Posted by Steve Waddington at 12:14

Saturday, December 8, 2007

PeerApp Results IV

The questions we had about the various aspect of performance have been answered by PeerApp. They are shipping us another Alteon switch which will solve the high CPU problem on the existing switch. Another of the 'puzzles' was solved when I realised I had been looking at the wrong graph for traffic loads - well not the wrong graph exactly, but the traffic on the wrong side of an interface - the traffic sent to the cache rather than the traffic directly from our supplier. This is the graph that connects to the Internet traffic supplier: This is the graph that shows the traffic redirected to the cache: And then this graph shows the net traffic to our border router: Having only been using MRTG for, oh, I don't know, twelve years, you would think I would know what I was looking at by now. Still nothing one of those 'd'oh' moments and a slap on the forehead couldn't fix. The important figure - the traffic generated by the cache, is fortunately not open to such misinterpretation. Last night it peaked at 260Mbps at midnight. The reason that we don't see higher cache generated traffic the rest of the time is, according to PeerApp, because only one external link is going via the cache at the moment.

The asymmetric traffic path is preventing the full effects of the cache to be realized. That makes sense to me. We expect to be able to add our second external feed on Monday. Actually, we have been trying for the last three weeks to do that, but have been held up by the suppliers inability to make the basic changes needed. I have been running ISP networks for a while now, having built what became two of the largest retail ISP networks in Australia, and built out and managed one of the largest wholesale networks. So I like to think not many mysteries remain to me when it comes to BGP peering. So when a large Telco supplier tells me that multi-hop BGP is 'complex' and 'non standard' and hence very difficult to do, I really have to wonder just who the hell they have working for them claiming to be network engineers, that they let loose on their production equipment. It has been a comedy of errors so far. First the wrong peer address was used. Apparently the person making the change must have typed it in manually rather than cut and paste from the email we sent. Not to cut and paste from the customers actual request is always a bad plan. The second time, can you believe it, they used their own interface IP address as the static route to our bgp loopback address instead of our router interface address. I don't even know how they could make such a mistake. Cisco IOS throws an error when you try to add a static route for an interface on the same router. I wouldn't have believed it in fact, except that they send Vinna the code they were using, and sure enough, it was their interface address, not ours. The supplier engineer must have just ignored the error message when wondering why it didn't work. All we are asking for is 1. change the peer address 2. add a single static route 3. add the line ebgp multi-hop to the bgp code. If, as a network engineer, you can't do that correctly in five minutes, let alone three weeks, well, there is a big question about your future in that line of work if you ask me. Anyway, hopefully, we have corrected all the things they can get wrong and we will have the change complete on Monday morning. In the meantime, I have thought of a way to give full symmetry to some of our network blocks via the cache. I think I will try that out on Sunday afternoon to give us a baseline on what to expect when both suppliers are redirected to the cache. Our middle child is sailing on the Leeuwin II today. Which has nothing at all to do with p2p caching, but is much more pleasant way to spend a balmy Perth Saturday than worrying about network performance. And a far more interesting graphic than MRTG graphs.

Posted by Steve Waddington at 13:27

Thursday, December 6, 2007

Interent Quiz

14 years on, and I have at least some grasp of the basics:

On the other hand:

95% Geek

Posted by Steve Waddington at 13:03

Tuesday, December 4, 2007

Servers as 'Bad' as an SUV?

In this New Scientist article, a point is made that a single server has the same carbon footprint as a fuel-inefficient SUV. I had read that report in some other publication a week or so ago, did a basic mental calculation, and determined it was an odd claim to make, given a moment to work it out. But since the article was fairly obscure, didn't give it any more thought. I am very surprised to see it in New Scientist - a publication I read weekly and have always had the highest respect for. Not the least being the editorial staffs numeracy and veracity. Sure, they are quoting another source. But if you are going to quote a source as 'true', surely some check should be made that is actually is. A single server has a similar carbon footprint to an SUV. Really? How? There are just so many ways it can't possibly be even close. For a start, any modern vehicle has about the same CPU power in various on-board computers and circuitry as a server would have, from the point of view of energy to manufacture and power to run. But I think the chassis the SUV computers come in is a little larger. My initial mental calculation however was based on the power output, and is easy to work out. The generic SUV referred to in the article is no doubt a US model, which is reasonable to assume would generate about 200KW - say 50KW average in normal driving circumstances. Normal use could be estimated at 2 hours per day - a average commute in most cases. It would therefore consume $50 \times 2 \times 365 = 36,500$ KW hours per year. A server, say a really big one with lots of hard drives, would have a 450W power supply, but actually draw, say, 300W. Obviously it would be powered 24 hours a day. Which gives $300 \times 24 \times 365 = 2,628$ KW hours per year. If they said fourteen big, powerful servers running 24 hours a day, have the same footprint as an SUV, yes, I could see that being the case. But not one. And the price of New Scientist has gone from \$5.95 to \$9.95 in the last year. Maybe they should use some of that money to buy a calculator. Or maybe they didn't want to use one because they figure the carbon footprint is the same as a Corolla. Incidentally, Exetel's server and router power requirements are currently 67,452 KW hours per year, or just over 1 KW hour per year per customer. Based on the, fairly difficult to interpret, 'green power' and carbon neutrality guidelines, the cost to an Exetel customer for a carbon neutral Internet service would be about \$1 per month. Now that I think about it, it would be interesting to see if that would be an option people would consider at all. I'll raise it with John.

Posted by Steve Waddington at 11:50

Saturday, December 1, 2007

PeerApp Results III

Just after midnight the cache hit 221Mbps of generated traffic. Again, the bottleneck was the, currently, single GE circuit connecting it to our border routers. I expect we will hear an answer about that from PeerApp on Monday. It should just be a matter of connecting a second GE port from the PeerApp switch to our border router LAN to give us 2 x 1Gbps capacity, but we need the PeerApp techs to set up their side first. The cache disk utilization is now at 20%, so it looks like the estimate of a week to 'fill it up' is about right. However one of the four redirector CPU's is reaching 90% and two others are peaking above 80%. That means that already we are very near max. capacity as far as the amount of traffic we are offering goes. Also, I think we will hear from PeerApp about what to do about that on Monday - it seems to me like a trivial problem, and nothing that a bigger, or second, Alteon switch, or a couple more server blades, couldn't fix. There are some puzzling things about the cache generated traffic though. I am 100% sure that the generated traffic is 'real'; that is, not just some measuring anomaly. I can tell this because apart from the PeerApp management interface, I am running simple MRTG graphs measuring traffic from the Internet and traffic from the cache. By taking the outbound traffic from the Cache and subtracting it from the inbound traffic from the Internet, I get a net difference of traffic, which is that generated by the cache itself. This also matches with 1% the value measured by the PeerApp reporting. In addition, there is no doubt my own first hand testing, testing by staff and feedback from customers shows a definite 4-5 times increase in P2P download speed as a result of the cache. Daily Graph (5 Minute Average)

The two puzzling things are: 1. The increase in traffic should be measurable on the LNS links and carrier cross-connects to our ADSL customers, but there is no evidence of that. 2. Why the traffic has such a sudden peak at midnight, and why it is comparatively so much lower than that later during the off peak period, when at the same time P2P traffic increases. However, there has been a measurable effect on the Verizon bandwidth, reducing it from a peak of 824Mbps to, last night, 587Mbps. That equates to a saving of 215Mbps on external bandwidth. A very promising result this early into the trial. My thinking about the above two points is that for 1. the increase is going to be spread out across several carrier cross-connects, making the increase less obvious than if it were on a single measurable circuit. We will be able to tell if this is the case because we should see a general increase in the combined graphs over a longer period of time. For 2. it could be due to some sort of interaction between the NetEnforcer and the PeerApp equipment (although the PeerApp has a 'straight through' path via the NetEnforcer). Or it could have something to do with the much faster download speed allowing popular downloads to be delivered quickly, then the generated traffic tails off leaving the more obscure requests that are less likely to be cached. It is still early days, and no doubt the interpretation and analysis is going to take a bit more practice to get right.

Posted by Steve Waddington at 13:27

Friday, November 30, 2007

PeerApp Results II

Last night the PeerApp generated 200Mbps extra traffic from 400Mbps being sent to it. Disk storage was 150Gbytes of 3900Gbytes total, or 3.8%. And the average CPU utilization for the four redirector processors was 62%. Overnight the disks have now reached 8.5% capacity, but after peaking at 200Mbps, the generated traffic dropped back over the next three hours to 50Mbps average for the rest of the night. This graph shows the P2P traffic analysis. The lower half is traffic coming into the cache from the Internet, the top half is traffic going to Exetel end users. The black line is the traffic delivered directly from the cache. Based on these current, early, figures, I make the following estimates:- the hardware supplied will be able to handle a total of 600Mbps redirected traffic, or not more than 900Mbps total traffic.- At 600Mbps with the cache fully populated, we should get 400Mbps generated from the cache. Of course, that is just with the hardware installed initially, for our trial period. The folks at PeerApp will assess that and put it into their plans of what will be needed in the long run. My own direct experience is that without the cache I would typically see 30-40kBps P2P downloads from heavily seeded torrents such as here <http://linuxtracker.org/index.php?sort=8> for example. With the cache, I am seeing 120-140kBps consistently - about as much as my 1.5Mbps ADSL line can handle. One bottleneck, which is obvious when you think about it, is we have hit the limit of the GE circuit between us, the cache and Verizon. 800Mbps from Verizon + 200Mbps from the cache = all we can get from that circuit. So the next immediate task is to load balance the cache traffic to the Exetel network over several GE circuits.

Posted by Steve Waddington at 10:33

Thursday, November 29, 2007

PeerApp Preliminary Results

It is too early to get any real results from the caching, but some early tests from a 1.5Mbps ADSL1 circuit are showing:
Typical d/l speed uncached - 40kBps
Typical d/l speed cached - 100kBps
We are currently offering only 120Mbps of traffic for possible selection by the cache redirector. Of that, 40Mbps is being classified as 'interesting' p2p traffic. That is a lower figure than I was expecting, since deep packet inspection tells us that over 50% of traffic is p2p. It may be that encryption and other stealth methods that people have used to try to fool the DPI of the NetEnforcer are precluding that traffic from selection by the PeerApp. After 24 hours of operation, the unit has performed without any disruption to other traffic and cache size has reached 17Gbytes, or 0.43%. There is still a ways to go before we can draw any real conclusions. We will be increasing the traffic we offer to the cache redirector to 700Mbps later today. I expect that to make a difference as well.

Posted by Steve Waddington at 13:30

Wednesday, November 28, 2007

PeerApp Setup

We have completed all the physical aspects of the installation of our trial of the PeerApp equipment. The process typified the, sensible, approach of 'plan for the best but expect the worst' (or is it the other way around?). Anyway, our plan was to thoroughly prepare and take as much time as needed to make absolutely sure we could do the needed changes with minimal disruption, and to have a fast roll back plan at any stage if something unexpected happened. As it turned out it went about as well as it could have. We changed our two Internet supplier circuits, 1.4Gbps of bandwidth, with less than 1 minute disruption to service. We moved half of our switching and routing equipment, and half a rack of servers, also with almost no impact to services. And then totally re-routed inbound and outbound traffic and changed the border network topology route-map policies with less than 30 minutes service impact during a planned maintenance window. A great job by the team; Ahad in particular, and Vinna also whose first time it was working on a change project of this scope. The PeerApp equipment itself is impressive; comprising of a Dell Blade server for the main processing, 4.5 Tbytes of EMC RAID fibre channel storage, an Alteon L2-7 switch and two Cisco 3560G gigabit switch route processors. Assembled, the whole lot only occupies 22 RU, but draws 19 Amps, so you can get an idea of the processor and storage density involved. The next step is to divert some traffic to the cache and test the results against what we expect. The network is only lightly partitioned in terms of IP address blocks, because for the first three years all of our customer IP's terminated at the Sydney POP and there was no need for subnetted groups. We do have a /18 block for the Melbourne POP that currently carries about 200Mbps of traffic, and also the staff IP address /24 block, as well as another /18 for Brisbane that carries 75Mbps at the moment. In combination, these will allow us a full range of test and comparisons to be made before we commit our entire /16 and /17 supernets. Testing the cache will start today or tomorrow, once the PeerApp folk verify the configuration required. What I would like to see is a 25-30% improvement on external traffic. It is expected the cache will take about a week to fully populate before it is at its most effective, so time will tell. We are all very interested to see the results.

Posted by Steve Waddington at 14:05

Monday, November 26, 2007

There's got to be a morning after...

Well 53% of voting Australians, I hope you were careful about what you wished for, because now you have it. Vinna (from Shanghai) mentioned on Thursday that she thought a lot of people of Chinese background would vote for labour because Rudd spoke Chinese. Maybe that did indeed add to the labour victory, who knows. For sure no one who speaks Mandarin can be a bad leader, or repress their people in any way. Maybe we can replace the promised 'education revolution' with 'cultural revolution'? One small step for Rudd, once Great Leap Forward for labour-kind? It is said it is better to be a dog in a peaceful time than a man in a chaotic time, but personally, I would rather be a man in a peaceful time. Unfortunately, it doesn't look like there will be much chance of that. We can certainly look forward to interesting times, whatever the case. ... if we can hold on through the night.

Posted by Steve Waddington at 09:27

Thursday, November 22, 2007

The Sky Is Falling II

The world, the IPv4 world of the Internet that is, is still turning. How unexpected. I think the human brain, in general, must have a problem coping with non-linear curves. For example, everyone can look at a straight line on a graph and extrapolate forward to some future value. However any curve in the line, such as geometric growth or the inverse of that seems to befuddle our predictive ability. I think it must be like a curve ball or a spin bowler. Any moderately competent batsman can tonk ball after ball to the boundary bowled straight down the wicket. It's the movement off the seam that is the problem. Maybe that is why the prophets of doom have predicted 'the end of all IPv4 address space is imminent' for the last ten years. If you take the tangent of the growth curve of IPv4 address space usage anytime during the late 90's, it was just about vertical, hence, in the 'very near' future, address space was bound to run out. So too for every year into the 21st century so far, some sort of straight line 'fit' to the curve shows 'very soon' all the address space will be used - the current 'very soon' being 2010 - 2011. Because 'historically' the growth has been phenomenal. And of course there is no reason to expect that what has happened before wont keep happening forever into the future, right? Keep banging the rocks together guys. As sure as Kilimanjaro rises like Olympus above the Serengeti, change is inevitable. So why does no one think the use of IPv4 will not change? The end has been predicted for a decade and never happened. As it stands, the total global utilisation of IPv4 is 60% - about the same as it was in 1999, though my memory of the exact figure could be in error. The reason the two are so similar is because of two things. First, growth of IP address use slowed because of the dot com crash. Second, previously reserved address space was released into the pool available for use. So the pool got bigger and the take up rate of the pool slowed. However in the last few years the take up rate has picked up again, and so the curve again is starting to look as vertical as it was in 98/99. Leading no doubt to the end is high predictions. The huge increase in utilisation this time has been fueled by DSL - where every connection requires a public IP address. But as everyone who can get DSL, gets DSL, the rate of use of IP addresses will slow again. Which seems to be the case as this article, about how demand for new residential Internet circuits is trailing off in the US, indicates. If the same isn't true now in Australia and other countries where Internet is widely used, it soon will be. Another factor is servers. There can be only so many servers fitting into so much room in so many datacentres in the world. I can find no stats to back it, but surely the number of server IP addresses needed must be starting to flatten out as well. Sure, manufacturers are producing ever more servers - but are they devices that need new IP addresses, or are they replacements for older devices and will take over those already used IP addresses? The answer is certainly 'both', but more towards the latter I feel. Exhaustion by 2010? I don't think so. It will be impossible to really find out, because IPv6 is going to happen no matter what, and I would guess maybe half the worlds current IPv4 addresses will be on IPv6 by then anyway. If it didn't happen, and I were to speculate, I would say IPv4 would be good to 2015 at the minimum and most likely still have life at 2020.

Posted by Steve Waddington at 16:55

Monday, November 12, 2007

Who's Time is it Anyway?

Today for the first time since I can remember I am taking a flight in the middle of the working day. My usual schedule is to take the Sunday midday flight to Sydney and the Friday night or Saturday morning flight back to Perth. This time however I have a stopover in Melbourne to do some work at the POP there, and the availability of flights to match when the equipment would arrive was to take the 1pm flight from Perth. I figured to find the airport fairly empty, and the Qantas business lounge all but deserted. I mean, who, apart from those stuck with a midday schedule like me, or people on holidays, would choose to take that much time out of their day? I got it so wrong. Many, many people it seems. The business lounge was full to overflowing and the gate lounges were just as full. Just who are all these people - regular travellers clearly - that can blow off a whole day in travel? A not so quick scan of the lounge (while looking in vain for an empty seat) showed: a) every second person had a laptop open b) every third person was giving the 'Solidarity Siblings' salute with their mobile phone encircling fists stuck to the side of their head. Well, I thought, at least if they are traveling on company time, they can use their laptops to do something productive. That would have been the case, except from the few screens I glanced at I am not sure how much solving Freecell solitaire games adds to the productivity of a company. At least the mobile phone call back to the office is productive, right? For sure - I think the reasoning goes like this: "I am out of the office in the middle of the day so I will call anyone I can think of back at the office to first make myself look important, and second leave no doubt in anyone's mind I am still being really productive travelling in the middle of the day". You know, if I am traveling during business hours, I see it as common courtesy to let the people who have to cover for me get on with their jobs without calling them to rub in the fact they have to do the work I have left. (another reason I travel on my own time whenever possible). So just who were all these people taking every available seat and hanging round the bar three deep waiting for the free drinks to arrive? Of course there were the smattering of suited travelers and 'grey pilgrims', and I guess the couple with the small kids were on holidays. The majority however, judging by the short sleeved and polo shirts, many with company logos, were involved in the mining industry. Hardly surprising in 'China's quarry' as I heard someone call WA the other day. Sure enough, a Sydney flight was called and no one moved. The Parapadu and Ravensthorp flights were called next and, hey, free seats all of a sudden. How many litres of natural gas will the North West Shelf Venture need to sell to China at 0.000012 cents per litre to make up for their eight or ten people traveling on Monday? No idea. A lot though. Good thing there is a lot of gas there. Make the most of it while it lasts fellas.

Posted by Steve Waddington at 13:22

Sunday, November 11, 2007

2nd POP Build Out

I have completed the first stage design to split the Exetel Sydney network across two POP's. It is something I have been reluctant to do for a long time, not the least reason being the extra cost incurred by the necessity of linking both sites with a fairly hefty piece of bandwidth. When the need for geographical diversity of POP's in a single city is analysed, it is pretty easy to see that no real technical argument holds up. It _used_ to, yes, six or seven years ago, when bandwidth utilisation were much smaller, the long term reliability of data centres, high speed devices and bandwidths was not as well known. But with seven years of extra data, those considerations that were unknown back in the day, and therefore had to be factored in as 'worst case', now have much better statistical credibility to back their 'five nines' claims. Two POP's basically means full duplication of everything in two locations so that should some unthinkable event happen and the connection between them fail, each can run autonomously, serving email, provisioning domain name resolution, and most importantly, continue to bill correctly, for those customers connected to each site. Surprisingly, apart from the duplication of bandwidth, the duplication of services does not add that much extra expense - mainly because almost all services are duplicated or have a backup server anyway. The 'cost' becomes the time and planning that needs to go into the logistical exercise of moving those services, without disruption, from one site to another. But as I said, such a move is in fact quite unnecessary if you look at the difference downtime expectation of one POP versus two POP's. Given the 30 minutes of downtime we have had at our single site (as a result of a UPS power failure) in four years - is it worth the extra \$15k per month to prevent, perhaps, 30 minutes of downtime in the next four years? Well I guess if you knew that 30 minutes would occur when you are on the final table of a million dollar on line poker tournament, your answer would be 'SPEND THE MONEY'. Otherwise, probably not. So why have two POP's then? Say someone was holding something of personal value to you to ransom or threatening you in some way. In my case, and I imagine most people are like this, you would not give in to their demands and tell them to do their worst. However, what if it was the welfare of your family or friends that were at stake? That makes it much harder not to give in. In fact, if the equation works out that the only way to ensure their continued well being (albeit slightly poorer well being) is to swallow some pride and give in to the bully, then that is what you have to do. So when a carrier insists that the only way to increase bandwidth past a certain amount is to have a second POP, of course you ask 'why'? The initial answer was something like 'because we say so'. Not a good enough reason, in my opinion, which I pointed out and over way too many iterations, kept asking to provide a reason or install the circuit. Well, finally, it turns out the reason a 2nd circuit can not be installed at the same location is because "It will confuse the field technicians if there is a fault". Ah yes, how obvious, because a carrier only ever installs one circuit in total per building in Australia - because two circuits, being more than one, makes it far too confusing to track down faults. Our options were 1) keep arguing for a single site, 2) take it to court, 3) accept it as inevitable. It takes a lot of stamina to argue with a large carrier - they can soak up blow after blow with seemingly no effect for a very long time. Eventually beatable, but in the meantime, everything else can fall to ruin. It may even be a deliberate strategy they use, since it is so often successful. Similarly for legal action. Years to get anywhere, and lots of money besides. We have thirty seven people with families, school fees, mortgages and credit card payments who depend Exetel for those things. And sixty five thousand customers who pay all the salaries and bills. There is only one choice. But man I hate giving in to bullies.

Posted by Steve Waddington at 14:11

Saturday, November 10. 2007

Razor

I believe Battlestar Galactica: Razor has recently made it's way to bittorrent. And for the last few days, bandwidth utilisation has been uncharacteristically high for the time of the month. Can it be that our little business, (to which I have devoted the last four years of my life, while my sons have grown up wondering where their dad is on the weekend, my wife gets invited to single mom's morning tea and I have been slowly nurturing with stress and worry what I am sure will turn out to be three different types of cancer) is nothing but a conduit for thieves? Surely not. Because that would mean if people say "Why are my downloads so slow?", "Why aren't I getting what I paid for?", "I'm going to the TIO, to report you for the rip-off merchants you are", it is the cry the thieves, disgruntled that the property they are stealing does not download at the maximum possible rate. The TIO is of course completely 'neutral', no way could you say that they are driven by revenue quotas or do anything other than to assess each 'complaint' on its merits with fairness to all and favour to none. No way are they the protector of every petty thief, serial nuisance and pathetic complainer that has nothing better to do with their time than whine about how unfair their lot is. What a depressing thought. Maybe I should do something to cheer myself up. I know - Why don't you all just *\$%& off you *\$%&ing thieving ^\$+@s!! Hmm, a small lift, but no real help. Oh well, the sun is shining, the peculator is perking and I won a big pot with an outrageous bluff at poker last night. Only another sixteen hours of work and I can sleep again. Ah, sleep. See, there is something to look forward to in every day.

Posted by Steve Waddington at 10:35

Thursday, November 8, 2007

Who's on First?

You gotta laugh. No, I mean really, you have to. Either that or tear your hair out in frustration at the total.. total... total frustration of it all. You see, dealing with large carriers is rather like being Costello when he just wants to know the name of the first baseman on Abbots baseball team. Every answer he gets is 100% correct, they are both communicating in the same language, what he is told seems to make sense, except for some 'piece' he seems to be missing that keeps him from really understanding what is being said. I kid you not, the conversation goes like this (over many days, but I have condensed it into a continuous dialog here):

Me: Can you tell me the lead time for the circuit we ordered?
Carrier: The circuit will be installed by twenty days
Me: Will there be any delays?
Carrier: Unlikely, no, it doesn't seem that way
10 Days later....
Me: Just checking to make sure the circuit is on track
Carrier: Yes, it will be installed by Twenty days as we said before
Me: So there are no delays?
Carrier: As we said, delays are Unlikely, no, it doesn't seem that way
5 days after that...
Me: How is our circuit going?
Carrier: Yep, still by twenty days
Me: Nothing is holding it up then?
Carrier: Unlikely, no, it doesn't seem that way
1 day before installation is due...
Me: Just checking that the installation is all go
Carrier: We have told you many times
Me: Is there anything that will stop it going ahead tomorrow?
Carrier: Unlikely, no, it doesn't seem that way
On the day....
Carrier: Me: Er, hello, you said the circuit was going to be installed...
Carrier: Oh, hello, yes, that's right, by twenty days
Me: So there is a delay?
Carrier: Unlikely, no, it doesn't seem that way
Me: So it is being installed?
Carrier: That's right, by twenty days
One or more days go by:
Me: So there has been a delay?
Carrier: As we said; Unlikely, no, it doesn't seem that way
Me: But you said twenty days
Carrier: Yes, the circuit will be installed by twenty days
Me: So what's the delay (it is at this point I am sure my sanity begins to break down, because I am sure I hear them say the next line)
Carrier: What's on second (no, that can't be right, I have misheard)
Me: Ok, let me make sure I understand this; There is no delay
Carrier: Precisely, no delay, also unlikely and it doesn't seem that way
Me: But you said the circuit will be installed by twenty days
Carrier: Exactly, as we have said
Me: So the delay is unlikely, there is no delay and it doesn't seem that way, and the circuit is still going to be installed by twenty days?
Carrier: You've hit the nail on the head
Me: Ok, so tell me who will install the circuit?
Carrier: No, the circuit will be installed by twenty days
Me: I know that, but by twenty days from when?
Carrier: Tomorrow
Me: The circuit will be installed by tomorrow?
Carrier: Tomorrow in provisioning, however the circuit will be installed by twenty days
Me: Provisioning will tell me tomorrow?
Carrier: If you ask tomorrow in provisioning, that's right
Me: Why can't you tell me today?
Carrier: I can't tell you today, today is out
Me: So you are saying that the circuit is still going to be installed by twenty days, but I can only be told by tomorrow because today is out and the delay is unlikely, no it doesn't seem that way?
Carrier: Exactly, you have it in one.
Me: But I don't have anything!! How can I get this sorted out
Carrier:Â No one can sort it out
Me:Â Surely someone can, who is in charge?
Carrier:Â I'm telling you, there is someone but they are away, no one can sort it out
Me:Â Well if there is someone, who do they report to?
Carrier:Â They report to nobody
Me:Â So you are saying that nobody is the manager, and no one can help?
Carrier:Â Can nobody help at all?
Carrier:Â Yes, nobody and no one can help
Me:Â So the circuit will be installed by twenty days, the delay is unlikely, no, it doesn't seem that way, you can't tell me when, but if I ask provisioning, they will tell me tomorrow, I can't ask today because today is out, there is someone who can sort it out but they are away, so there is no one who can help, and nobody is in charge?
Carrier:Â Precisely.Â My, you are a bit of a guru at this aren't you.Â Not many people understand it as clearly as you do.
Me: Please, please, just tell me who can sort this out?
Carrier: I can't tell you who can sort this out.
Me: Why not????
Carrier: Who's on first
Of course, it isn't really like that. I have left out all the non-returned calls, emails not replied to until the third or fourth followup, no one available due to RDO's etc, etc, that make it MUCH more frustrating. What I can say, is that in over twenty years of ordering data circuits, which would number in the hundreds in total, not one, ever, has been installed by twenty days.

Posted by Steve Waddington at 11:38

Tuesday, November 6, 2007

New additions

From time to time we have had a people ask for services such as BGP, QoS, MPLS and being able to co-locate equipment at the Exetel POP. The problem with those sort of requests are that, firstly, very few people want, or care about them - in fact the requests are usually limited to about one or two a year on average. So that to offer them as part of a 'standard' product makes the product more complicated and harder to understand for those that don't care, and creates more complexity in the network for the benefit of only a very few. Secondly, of those that do make the request, most often it is because they picked up on the buzz words, but haven't really planned through what it will mean to them. Here are some examples of how the end user planning fails to make it to the first, obvious, question: BGP - Do you have an assigned AS number, or do you plan to use a private one? Most often the answer comes back "I need a what now?" QoS - the end user question is "Will your network support QoS end to end", to which there are many responses, but usually I go for "As far as the Exetel network goes, it is run without contention. Do you mean you want to run QoS on your WAN links to control your own traffic, or somehow have your QoS policy extended across the global Internet?", but said more politely of course. MPLS - Do you plan to use BGP, OSPF or some other protocol for your VRF, and who is your network engineer who will manage your route policies? Co-Location - Have you considered access to your equipment, remote management, monitoring, how alarms will be attended to, and the inevitable delay and cost for on-site attendance? What really bugs me is that more often than not, these requests come from IT consultants acting on behalf of their 'client', and who have put enough thought into the solution to come up with the buzz words, but exactly no more than that. Anyway, I've said plenty on that subject previously. Never the less, there is from time to time a request for those extra services where the need is real and based on a real understanding of what is being asked for. Two such cases have come up just recently, one for BGP peering and one for co-location. For BGP, the customer's contract engineer was a person I worked with many years ago, who was competent and sensible then, and could only have gained more experience since. So on Friday when she asked if we did BGP to allow fail-over of a Metro ethernet link to SHDSL, I explained we didn't but would be happy to work with her on a solution. Which we did. I thought it through for a few days, ran the solution past her, and by the following Friday we had our first customer BGP session up. It is still the case that very few people will want or need BGP - but it is a nice, fully engineered, addition to the service range that is now available. Co-location, as a technical product, is very straight forward - bolt some equipment into a rack, bung in an ethernet cable, switch on the power and away you go. Well sure. But logistically, there is far more to consider. So much so that the logistical issues often outweigh any benefit of position, especially considering today's very high speed and relatively low cost data connections available to just about any business location. But again, when a customer comes forward with a request based on a legitimate need, has an understanding the costs and given full consideration to the logistics, it makes for a good case to see if they can be accommodated. And so too will we shortly commission new racks in the Pymont and CBD data centres for the purpose of dedicated hosted servers. So, those two new product additions will help further 'round out' our business services range. And since these things always come in threes, I guess I will anticipate the inevitable request and begin working on a QoS solution we can offer for the CPE.

Posted by Steve Waddington at 11:10

Monday, November 5, 2007

Eliminate Spam II

Call me slow, but a new form of spam has reared its ugly head - bogus tracebacks on blogs that lead to spam web sites. I have just spent the last 20 minutes deleting a mass of tracebacks added to one of my blog posts overnight. Needless to say, the next thing I did was look for the 'Disable tracebacks' option in the blog setup - but could not switch it off in quickly enough to prevent another three spam tracebacks being sent. Once one of those botnets gets hold of anything remotely spam exploitable, there is no letting go. The traceback links were many and varied, and mostly unrepeatable on a site where there is some chance children may see them. But a bit of work with dig shows them all to use the two name servers 62.213.64.1 and 62.213.64.192. And all the sites resolved to an IP address in the range 62.213.64.0/24. Whois shows the sites were mostly .cn TDL's, but the IP address range is in Russia: inetnum: 62.213.64.0 - 62.213.64.255 netname: S-TELEPORT-NET descr: Sputnik Teleport Ltd. descr: Moscow Holding Company country: RU admin-c: AAM5-RIPE tech-c: AAM5-RIPE status: ASSIGNED PA mnt-by: CARAVAN53-MNT notify: pez@obl.ru changed: skiv@caravan.ru 20011122 source: RIPE It is hard to say if it is a legitimate network that has been exploited, or a dedicated spam network. Whatever the case, having direct access to border routers means positive action can be taken: #ip route 62.213.64.0 255.255.255.0 null0 No more spam from that network. I don't think spammers like people who take action against spam, because out of all the blog posts, the only one that attracted the spam tracebacks was the one called "Eliminate Spam". Ah well, I'd like to see them get around a null0 route!

Posted by Steve Waddington at 10:21

Saturday, October 27, 2007

I Smell A Rat

I have seen the BigPond ad 'those homes, those homes..' a couple of times now, and since I first saw it, something struck me as odd. The family portrayed are obviously fairly affluent, in their three story, nicely appointed house (wooden house mind you. Looks more American than Australian to me. Maybe the ad execs were playing the 'no place like home' card). Yet there is a big rat in the ceiling between the third and second story. Huh? It's not a mouse, or a opossum, it's a big fat rat. Now I don't mind rodents, and I have no phobia towards them. But I don't like them in my house. Many people I am sure have far greater negative reactions to them than I. My wife for one. Let's face it, the best possible reaction to a house with rats is dislike, and it only gets worse from there, down to the utter revulsion no few people would be bound to feel. So why have a rat there at all? No one in their right mind puts a big dirty rat in the bracket of 'cute furry animal to soften the home' scene. Yet no doubt hundreds of thousands of dollars was spent on the ad, and millions on the campaign. So what is it doing there? There is no way it could be there by chance. It has to have some meaning. Well, I obviously have no idea what must have been going through the heads of the people that decided to put it there. But I do like those little clues to meaning that artists like to put in their paintings, particularly portraits. You know, where the subjects finger points towards and open book that is on a page that has some relevance to their life, or the colour and curve of the drapes in the background is symbolic of something of the time period. Maybe so too does the rat have some particular meaning in the ad. Could it be it is a warning? Is it saying "on the surface, this product we are selling looks very nice, but beware, there is a rat hidden here"? So who would do that? It seems unlikely Bigpond would actually try to imply such a thing of its own product. Any hidden costs or potential disadvantage to the consumers would have to be fairly disclosed, surely. There would be no question a company as meticulously ethical as BigPond has to be would make any disclaimer as completely public and obvious as possible. Yet there it is, a big, fat, dirty, rat, in the ceiling space of that very nice family's home, right by where the telephone cables would run. Since I have already moved well beyond the realm of the vaguely plausible to go any further and offer any sort of reason would be just pure conspiracy theory. Why is the rat there? I don't know. If anyone has a plausible explanation that would make any sense, I for one would like to hear it.

Posted by Steve Waddington at 13:44

Monday, October 22, 2007

Quiet day for New IPTV

A quiet day on the network front. As I predicted, the bandwidth variations stopped after a few hours - almost certainly when whatever upgrades or changes were completed on whatever network they were being done. It was also a good test of the load balancing script - it behaved as expected, shifting traffic from one supplier and back again as the load shifted. When it couldn't adjust any more, it diligently reported the fact that it had tried its best, but had run out of network blocks to move. We added a third party IPTV server as well today. Our plan is to make it available through Pipe peering only initially. It will give us a very good idea of the bandwidth profile and network impact of such a service.

Posted by Steve Waddington at 20:22

Sunday, October 21. 2007

Bandwidth all over the place

Something odd is happening with international bandwidth at the moment. Flows are changing by more than 100Mbps between our two suppliers, Verizon and Optus. I can see here that Global Crossing looks like have some interconnect issues, but that wouldn't affect us at all. I would say traffic on that network would account for less than 5% of inbound data to us. From the Internet Traffic Report it looks like there is a major disaster in progress, until you look closer and see two of the six relatively obscure networks they are getting data from are down. So the answer isn't there. This might be a clue; lots of packet loss at MAE-West, MCI (Verizon), Sprint and others. No reference to that data though, so it's hard to gauge its relevance. Time for some direct testing. Trace paths from us out all look ok. Pity there is no looking glass at MAE-West any more, but route server I have access to at PAIX seems to show things from there back to us are ok too. No surprise that nothing reported on the NANOG list, worth a check just in case though. I can't see any obvious cause. My guess would be that some major backbone in the US is shuffling traffic around late on a Saturday night (for them), possibly planned maintenance even. And other networks are manipulating route paths to compensate. Oh well, the ditch for the new pool light power line is dug. Three hours to the Rugby world cup replay. I expect it will settle down by then. Time to settle back with a book in one hand and an eye on the graphs in the other (yes I know, my hand-eye coordination was never good).

Posted by Steve Waddington at 13:27

Thursday, October 18, 2007

Jing's Fantastic Script

One of the Big costs in the operation of an ISP is Internet Bandwidth. I recall when Noah bought bandwidth for the Ark for \$20,000 per Mbps, and it doesn't seem that long ago to me - must be getting old. Today's prices are 1/50th or less of that, but even at that comparatively low price, it is not uncommon for an Internet end user have 75 to 100 or more times the bandwidth of a user of yesteryear (and use it). We now have over 1.6Gbps of Internet bandwidth, and one point six billion of anything, no matter how cheap, is still a lot of money. Right from the start we saw it as important to get the best, most use from our bandwidth. On a day by day and sometimes even hour by hour basis we have balanced the bandwidth usage across our Internet suppliers to maximise the use of the circuits. I have discussed in an earlier article how we have used BGP to shift data traffic load from one supplier to another, so we can balance the circuits to within 5Mbps each. That has always been a manual process, because a) there is no scripting or scheduling feature in Cisco IOS and b) BGP load balancing is more of an art than a science - a value judgement has to be made on what the curve of a graph means, and what do do about it. However, having done that task for several years, I recently handed it over to Jing, and in doing so, and having to explain the process in sufficient detail for someone else to understand, I began to see how it might be possible to automate that process all together using the NetSense scripting tools and a checking script with a few sanity controls. I talked with Jing about how such a script would need to:- Check the data rate on each border router for each supplier every five minutes- calculate a three sample average- check if the rate was above or below a preset limit for a supplier- if it were, add or remove a /24 block from one supplier to the other (shifting about 5Mbps of traffic)- have a web interface to maintain the preset values- apply sanity checks to the data and do nothing (but report) if the checks fail. I think Jing said something like 'that will be easy to do'. Well, AFAIK no one else in the world has ever been able to automate such a thing, but sure, I guess when you list them out, those six points _look_ easy to do. If it looks easy, then it must be easy, right? Though maybe 'easy' in the same way that an Olympic class slalom skier makes hurtling down a slope executing 6 g turns at breakneck speed look easy. But if not easy, then certainly do-able. So off Jing went and did it. For the last couple of days the script has been running, patiently adding or removing bandwidth at 5Mbps a time if a circuit need it (so most of the time doing nothing), each five minutes, 24 hours a day. The really pleasing thing about the way the script works is that the traffic graphs are the smoothest they have ever been. The micro-manipulation of traffic twenty times an hour is exactly what is needed to ensure the absolute optimum usage of our bandwidth. It's nice when something works out well like that.

Posted by Steve Waddington at 11:02

Wednesday, October 17, 2007

Carrier Licence

I've completed the technical part of our application for a Carrier licence today. At least I think I have. At first inspection it looks very involved. Lots of cross referencing to the Telecommunications Act, definitions of 'Network Units' under the act which must be named and numbered in detail. All running to hundreds of pages across multiple web sites and government publications. But then when it comes to the actual application itself the change is dramatic. There are two forms, both asking for essentially the same information, both only six pages. Two pages are for company details, address and signatures, two telling you what the fees and charges are. One page that is a checklist to make sure you have done everything, and one page that is the actual application itself. To that one page must be attached a description of the network and a diagram, along with identification details of each 'network unit'. Even the network unit turns out to be pretty straight forward (if I have understood it correctly). It is the part or parts of the network for which the Carrier licence is required, as defined by the TC Act. The definition of which takes about three or four paragraphs in the Act, and no problem to comprehend. I remember from past companies where the act of getting the Carriers Licence was the 'Big Thing' that occupied the time of a bevy of people for an inordinate amount of time. I just can't see what all the fuss was about now. Maybe I am the guy who finishes his 4 hour calculus exam with an hour and a half to go, only to realise after wasting the next 80 minutes that there is another side to the question sheet. Or maybe getting the licence isn't that hard after all. Ah well, hope springs eternal.

Posted by Steve Waddington at 05:42

Monday, October 8, 2007

Engineering a Network For VoIP

Is surprisingly simple, once the concept of how a voice signal as data must be treated, is understood. Which is 180 degrees away from how data as data would like to be treated. The differences are: Voice is intolerant of delay, most data protocols in IP can withstand 30 seconds or more delay with serious consequences. A delay of a several dozen milliseconds can be enough to make voice unintelligible. As well as a continuous stream of data, voice also likes little variance in the time each data packet arrives. Data protocols are oblivious to that need. Voice is very tolerant of packet loss, if only 80% of packets arrive, the signal is legible. Data requires 100% certainty of transmission and uses retransmission if necessary. Retransmission of lost voice packets would be a disaster. A lot of effort, at large router CPU cost, can be put into so called QoS controls for networks to overcome the apparent conflicts of carrying voice over a data network. But when what actually happens in voice and data transmission is looked at closely, it can be seen that such controls, except in narrowly defined circumstances, are mostly unnecessary. I can only agree with a comment Geoff Houston made at a conference in San Francisco some years ago, that "... after several years of trialing and using QoS, the conclusion we [Telstra] can come to is that if it is done carefully, it doesn't cause too much damage on the network". I think the need for QoS grew out of a time where data packets were large in relation to line speeds. For example, take a 1500 byte ethernet packet on a 128kbps serial line; the time to read that packet is in about 100ms. Line a few of those up in front of a 64 byte voice packet, and all of a sudden you have half a second or more of latency variance, which effectively makes the circuit unusable for voice as data. But I doubt there is any service provider today that uses anything less than 155Mbps, and most would use either 622Mbps, 1Gbps or multiples thereof in their network backbones. Even with jumbo packet sizes of 4kbytes, line serialization is a fraction of a millisecond. Unless the network was totally jammed with data, there would be no audible impact on the voice data transmission. Of course, if the network is indeed jammed with data, then something is needed to let the voice data packets through in a timely manner. And that is going to use lots of CPU cycles in whatever router is assigned the QoS task, because each packet has to be examined and matched against the QoS policy. However, we just want to get voice data packets through with, oh, say less than 10ms transmission variance, right? And we can stand a little bit of loss. (Consider though, that if we had no contention on the network in the first place, there would be no need for QoS (unless interface buffers are ridiculously large, which is never the case, so let's not worry about it). So having enough bandwidth capacity is a very good place to start before committing lots of router resources to QoS) Commercial and/or logistical, reality in Australia means that at times, some level of contention is inevitable. We should rush off to QoS controls then? No. Because the same thing can be achieved in any event through a sensible queuing policy on the router interfaces. Which sounds better to you? A queuing policy which says "small packets have inversely proportional priority so they don't get stuck behind big packets". Or a QoS policy which says "examine the contents of every packet and determine if it is between these sets of values, if it is, let it through up to this limit, if there is capacity to do that, otherwise send it to another policy to decide what to do with it, and if it isn't but the volume of them is below this threshold then send it through, but if the number is above this threshold, then put it in this container until it is full, but if the container is full, drop the oldest packet in it and store this one for a while". Hmm, tough choice. In fact, even a queuing policy isn't strictly necessary in most cases because of what I mentioned before about interface buffers. On a gigabit Ethernet interface, Cisco recommends an input hold queue size of 2,000. Even with a full queue of maximum size packets, the total time a packet can spend waiting for transmission is only 2.4 milliseconds - well within the tolerance level, by more than a factor of ten, of even the most finicky VoIP protocol. The real life test case is to see what happens when there is contention on a non-QoS's network and compare that with the alternative. The results are that even when bandwidth has so much contention that data packet loss (and voice packet loss too) is 5%, voice, due to its loss tolerant nature, is still understandable. With QoS, the packet loss for data kicks in earlier, so at 5% data packet loss, voice works fine. At 10% packet loss, the total network traffic is about the same as the 5% level on a non QoS network, and voice is still working. But which commercial network can stand 10% data packet loss? The company trying to flog that service 'because it had QoS' would go broke in a flat second. Not that 5% packet loss is good either, but for short periods of time, on an otherwise trouble free service, it can sometimes be forgiven if it is fixed very quickly. (oh come on, who am I kidding, it is never forgiven, it just has to be fixed, real fast, if it ever happens). It comes down to this; QoS in the service providers network core requires five times the expenditure on router processing power, which allows the service provider to offer usable VoIP when packet loss for data is 10%. Well, I should have no problem getting the CAPEX for that past the board. Though I rather think the need for QoS will be self correcting if the network is ever that lossy.

Posted by Steve Waddington at 15:32

Thursday, October 4, 2007

What is a Business Blog?

I received this almost-spam 'tip' from MelbourneIT (several times in fact) - <http://www.successbox.com.au/online-marketing/blogging-ups-your-web-ranking/Woo-hoo>. Very timely, I guess, if you otherwise have had no exposure to Internet trends for the last several years. I have to admit though, it did take me a while to 'get' what blogging was all about. As a business tool that is. And then a little while longer to do something about it. But there is no doubt the humble blog now has a cemented place in business life. Then I stumbled across a post on a forum from the General Manager of another service provider company who wrote: "And what is a "business blog"? Is that like Myspace for people in suits?" Yes indeed, ignorance is still alive and well in Perth. Nice to see our competitors with their heads firmly planted in the sand.

Posted by Steve Waddington at 13:27

Tuesday, September 25, 2007

The Job Isn't Done Until Your not in It

Some years ago, working at a startup that had serious aspirations to be a player in the global Telecommunications market (and still is in some countries), one of the Managing Directors made the comment to me "Steve, how would the network run if you weren't here". It wasn't said as a question, but I took it as one never the less. Because it is a question every systems architect and engineer should ask themselves about any project they work on. I have seen many people I have worked with - far too many in fact - despite whatever other fine qualities they possess, make themselves part of the running process of the systems they build. Some have noble reasons I am sure; with their intimate knowledge of the design, they are the best person to ensure smooth operation in production. Some I suspect have less noble reasons; if no one else knows how to run the system as well, they reason their tenure is assured. Others do it no doubt out of laziness, or that the effort of just getting it to work has taxed them to the extent that the further effort of documentation and operational processes is just too much for them. Whatever the reason, and despite how well the system may work, if the designer/builder is needed as an integral operational component then the system has a fundamental design flaw. At best it can be said to be still a prototype. The operational processes are as fundamental requirement of any design as the hardware the system runs on. The goal has to be that anyone with the basic skills necessary can pick up the operation of the particular systems. Inherent in the design must be designing the designer out as soon as the production stage is reached. Of course this applies to all aspects of, in my case, network engineering, including the running of the Network Engineering Department itself. The person that built it is usually not the best person to keep it running (unless they see it as their exit strategy to retirement), so hiring, nurturing and encouraging the right people should be of equal consideration. It is all part of the grand design, that takes me out of the equation, and completes the job. So my answer to the non-question was to blink and say 'Everything will just keep working because that's the way it was built', adding silently 'of course' at the end. And so it remains something I look back on with pride, that for several years after I left, until the receivers turned of the power, that network just kept working.

Posted by Steve Waddington at 10:56

Tuesday, September 18, 2007

next-hop-self

A difference between bgp and other routing protocols is that when bgp announces a network to a peer, it sends the next hop address to reach that network, which is always the next hop beyond the router itself. All other routing protocols simply announce a network and the peer sends traffic for that network to the router making the announcement - that is, the peer router implies that the announcing router is the next hop. The announcing router says "I have a route to this network", and the peer router sends traffic to that network by creating a route table entry like this: ip route [network] via [router IP address announcing the network]. However in bgp, the announcing router says "I have a route to this network, and here is the next hop that I use to reach it". The bgp peer router then adds the NEXT HOP to its routing table, like this: ip route [network] via [IP address I have been told is the next hop]. Which is just fine where the announcing router also is announcing a path to the next hop address, or the peer router has a valid route path to the next hop address. The problems next hop cause in eBGP are well known, and standard practice is to add next-hop-self for all eBGP peers. However the addition of the next hop information can cause a network to become blackholed in iBGP in some cases, as this example shows:

1. A subnet is announced by two IGP's, in this case iBGP and EIGRP.
2. iBGP announces the subnet to peers with the next-hop parameter, which is the WAN address to reach the subnet. EIGRP announces the subnet to peers with the local router peer address as the next hop.
3. Edge routers exchange routing information via EIGRP, edge to core is via iBGP.
4. The edge routers can reach the subnet via the next hop of the LAN ip address of the router announcing the subnet.
5. The core routers are told that the next hop to the subnet is the WAN ip address.
6. The core router however may not necessarily have a valid, or the correct path to the WAN IP address - this may be blocked by filter rules, passive interface rules, or propagated a different way by route policies. This effectively blackholes the traffic for that subnet.
7. By enabling the command next-hop-self on the edge router for its iBGP core peer, instead of sending the next hop it knows about as the path to the subnet, the router now sends its own IP address as the next hop.
8. The core routers now send traffic for the subnet directly to the edge routers, since that is now the next hop for that subnet. This resolves the problem.

Posted by Steve Waddington at 12:46

Monday, September 17, 2007

Foundry and Cisco

I have been using Foundry as core L2 switches since, hmmm, 1998 I am pretty sure. A long time anyway. An argument for not using non-Cisco equipment on a Cisco router network has always been the hassle of figuring out a problem where with of the two different manufacturers equipment could be the cause. The fear is that both will deny it is their own equipment at fault, and both will blame the other equipment and not investigate possible causes on their own equipment any further. Fortunately I hit on a method to solve that particular problem early on. No manufacturer, not even Cisco, can afford to be the 'lone wolf' of the industry. The implication that one set of standards based equipment wont interconnect with another set of standards based equipment is enough to motivate both TAC's into finding the real cause. But apart from a few extreme cases, that hasn't proved necessary, in all cases the Foundry switches have worked superbly with Cisco routers. Router to switch connectivity has never been a problem, even with the heavy use of dot1Q subinterfaces that are a hallmark of these later generation networks. The only problems have been with switch to switch connectivity. I have to say though, that for whatever problems I have run into with Foundry-Cisco switch connections, I think I have hit more with Cisco-Cisco switch connections, such as with 2900's to 3500's, 19xx/29xx switches to 45xx and any of those to 65xx models. Which isn't all that surprising when you consider they were all made by different manufacturers before Cisco bought them. For many years, the Foundry Edgellron 24G has been a favourite of mine. Excellent performance, and having deployed many dozen over the last five years and never having a single problem with any of them, I would have to say fairly reliable too. Unfortunately all good things must come to an end, and Foundry have retired that model and we can no longer buy them. The replacement model is the FLS6xx, which has a different OS, made to look more Cisco-like I think. And there in lies the problem. Make a switch CLI more like a Cisco switch is fine - but which Cisco switch exactly? Whereas our Foundry 24G interconnected seamlessly with our suppliers who have Cisco switches, we have now tried five times, and needed to roll back five times, getting the new switch to connect properly to a supplier switch. Three of the four suppliers who have Cisco switches have no problems, but one supplier with a Cisco 35xx switch we just can't get to work so far. The latest problem is cross-connect circuit comes up, and the L2 link is fine, but there is L3 packet loss. A switch causing L3 packet loss when there is no L2 problem? Bizarre. Well, there is going to be an answer. One way or another it will be resolved in the next few days, there is no doubt about that given the level we can escalate this to in respective TAC's. But I will lay odds of ten to one that the solution is going to be something simple.

Posted by Steve Waddington at 14:48

Friday, September 14, 2007

I want to go into Management

Almost everyone that I have hired has said that their goal at some point in their career is 'to go into Management'. Each one of those people also, at the time they make that statement, have no idea what that really means. I think what they understand a management position to be is a path to a higher salary, authority and seniority. And all that is of course true. What very, very few people understand at the start of their career (and sadly all too often in the middle and even end of their career), are the other, really hard things, that they have to do to a) become a manager and b) be effective as a manager. By 'Manager' I don't mean a title that is printed on a business card. As John pointed out to me the other day, instead a company that gives someone the title of 'Manager', could just as easily make the title 'Complete Idiot'. Labeling someone as a complete idiot no more makes them a complete idiot than labeling someone as a manager makes them a manager. It is only the individual's ability and personal characteristics that make one or the other true. If someone asks to be a manager, no organisation can somehow make them a manager, all that can be done is to provide them with the opportunity to become a manager. It is the individual who must decide if they are prepared to do the work and make the sacrifices required to turn the opportunity into career success.

Posted by Steve Waddington at 15:10

Thursday, September 13, 2007

We are here to help you

We had a surprise visit today from a major supplier. Not a surprise in that we didn't know the meeting was going to happen, it had been arranged a couple of weeks prior. Rather, a surprise in the value they brought to the meeting and genuine concern to cooperate and assist with our business growth. The agenda was purely technical and technical fault reporting issues. I have worked with this supplier at Exetel and other companies for over six years now, and they really have developed the reporting data from their network management systems to a high level of sophistication - much more so than their competitors who are also our suppliers. It is not something I had come to expect from any Telco supplier in the past. Many suppliers over many, many years have paid lip service to "we are in this together and we will work with you to sort out problems". Which invariably has meant "we are here to take your money and deny blame for any faults you encounter". But today, for the first time ever when a supplier, a Telco supplier has said that, I actually think they meant it, not only that, but were demonstrably capable of delivering on it. Did I mention they are a telco? I am clearly missing something very important that can only be detrimental to Exetel here. I am off to the doctor now to ask for a double strength prescription of cynicism pills.

Posted by Steve Waddington at 15:31

Wednesday, September 12, 2007

Controlling BGP traffic with Cisco IOS

1. A peering session to another BGP speaking router is initiated in the config router bgp section:
router bgp 10143
neighbor x.x.92.57 remote-as nnn
2. The filter-list command is used to control which ASN's the router will accept and send:
router bgp 10143
neighbor x.x.92.57 remote-as nnn
neighbor x.x.92.57 filter-list 40
in
neighbor x.x.92.57 filter-list 3 out
2.1 The filter-list command reads the corresponding ip as-path access-list:
ip as-path access-list 3 permit ^\$
ip as-path access-list 40 permit ^nnn_!ip as-path access lists use unix-like regular expressions to match ASN's. ^\$ matches the start and end of a line with nothing in it. ^nnn_ matches a line starting with nnn and anything following it.
3. The route-map command is used to apply one or more route-map rules to network addresses the router will accept or send.
In most cases, it is used to control specific IP address blocks:
router bgp 10143
neighbor x.x.92.57 remote-as nnn
neighbor x.x.92.57 remove-private-as
neighbor x.x.92.57 filter-list 40 in
neighbor x.x.92.57 filter-list 3 out
neighbor x.x.92.57 route-map bgp-full-default-in in
neighbor x.x.92.57 route-map bgp-peer1-out out
3.1 The route-map command reads the corresponding route-map rules in numerical order:
route-map bgp-peer1-out permit 10
match ip address prefix-list bgp-peer1-out
route-map bgp-peer1-out permit 20
match ip address prefix-list bgp-peer1-out-prepend1
set as-path prepend 10143!
route-map bgp-peer1-out permit 30
match ip address prefix-list bgp-peer1-out-prepend2
set as-path prepend 10143 10143!
if the route-map instruction is 'permit', it will allow any network permitted in the prefix-list to be announced by bgp.
If the instruction is 'deny' it will block any matching network permitted in the prefix-list.
3.2 The ip address prefix-list command reads the corresponding ip prefix-list of network addresses:
ip prefix-list bgp-peer1-out seq 1 permit x.x.128.0/17
ip prefix-list bgp-peer1-out seq 2 permit x.x.0.0/17
ip prefix-list bgp-peer1-out seq 3 permit x.x.0.0/17
ip prefix-list bgp-peer1-out seq 9 permit x.x.64.0/18
ip prefix-list bgp-peer1-out seq 19 permit x.x.0.0/24
ip prefix-list bgp-peer1-out seq 20 permit x.x.1.0/24
ip prefix-list bgp-peer1-out seq 21 permit x.x.5.0/24
ip prefix-list bgp-peer1-out seq 22 permit x.x.6.0/24
ip prefix-list bgp-peer1-out seq 23 permit x.x.7.0/24
ip prefix-list bgp-peer1-out seq 24 permit x.x.8.0/24
ip prefix-list bgp-peer1-out seq 100 permit x.x.236.0/24...
ip prefix-list bgp-peer1-out seq 1000 deny 0.0.0.0/0 le 32
The last line of this prefix-list will match any network of any netmask length.
In other words, any network not explicitly permitted, will be denied.
4. To summarize, in BGP:
* The filter-list command uses an ip as-path access-list to permit or deny AS numbers
* The route-map command uses a route-map apply route-map rules
* The route-map itself uses one or more ip address prefix-list to permit or deny network address blocks
IMPORTANT NOTE
All Cisco IOS access-lists and prefix-lists have an IMPLICIT 'deny any' rule at the end.
But more importantly, if you refer to a list and that list does not exist or has been removed, then the router will presume the implicit deny any rule.
This means that if you delete a list but the route-map or filter-list is still there, then you have told the router to apply 'deny any' for that list.

Posted by Steve Waddington at 14:05

Tuesday, September 11. 2007

Eliminate Spam

Well over 90% of all spam is sent by compromised PC's. Spam sent by dedicated spam servers are easy to blacklist. It should therefore be possible to eliminate spam world-wide by a) restricting Internet access of compromised PC's sending spam; b) forcing owners of said PC's to take anti-virus and anti exploitation measures and c) null routing the few remaining deliberate spam servers. We have been using such a method for several years now. While it does nothing to stop the flood of spam reaching Exetel, it ensures Exetel as an ISP has a far lower ratio of end users creating spam. Spam is reported to Exetel by other ISP's, in particular AOL has an excellent opt-in spam reporting systems for other ISP's to use. Almost all spam sent in the world makes it to at least some AOL users, and AOL reporting system has very high integrity. Of the many thousands of spam reports received, there is yet to be a false positive. The common format of the AOL, and some other ISP, spam reporting means the email notices can be parsed by an automatic processor. That processor strips the Exetel IP address and creates a list of IP address which are then sorted and run through the unix/linux utility 'uniq'. The result is a list of single IP addresses from which we can be certain spam has been sent. That list is then submitted to a database utility that matches the IP address to the end user. From there another utility changes the end user IP address from a public IP to a private IP address on a specific 10.x.x.x address block. The next step is to initiate a disconnect through radius so that the current user session is dropped and when reconnected, acquires the private IP address. All outbound traffic on the Exetel network is directed through core routers. Those core routers have specific route-map lists that look for 10.x.x.x addresses and redirect them to captive portal servers. The really nice thing about this method, is that it costs very little in terms of router CPU, because of the very specific nature of the traffic the route-maps are looking for. The result is that all identified sources of spam are: a) blocked from access to the global Internet, so they can do no more harm b) Directed to a captive portal web page that provides specific instructions on how to resolve their spam/virus issue c) except for the expected exceptions of those very ignorant people that can not understand written instructions, there is no intrusion of extra support burden placed on the Exetel helpdesk support as a result of spam. d) End users are forced to consider and take action on the inconvenience they cause to all other Internet users as a result of their carelessness. Really, it is a win for everyone. Side note: Who do you think are the most critical of this method? Of course it is the self fancied 'IT Gurus' who sell themselves as so-called experts. Yet despite their self acclaimed 'expert' status, are unable to prevent their own servers, or their long suffering clients, from sending spam - that the most basic of competent prevention methods would eliminate.

Posted by Steve Waddington at 17:09

Monday, September 10, 2007

The Best IGP

For a long time I maintained that EIGRP was an ideal IGP for any local Service Provider network. The reason I liked, and still like, it so much is that it is very simple to setup, for all intents and purposes has zero maintenance requirements, and above all, has very fast convergence. These days of residential services coming to expect close to 100% uptime, and the much, much faster connection process DSL modems that use PPP connections have, fast convergence means that Exetel's part of the service provision chain imposes no extra connection latency on the end user. In theory, if the Cisco design guides are followed, EIGRP has no limits of scale, except where the network routers themselves run out of processor CPU as a result of traffic. Well, that's the theory anyway. In practice, I found the limit at a network terminating 50,000 IP addresses. Which I have to admit, in olden days (aka dial-up) terms is massively bigger than anyone would every expect a local network to be. In the past of course, even if there were 50,000 subnets at a single site, the site itself would be broken into many smaller local networks, and the EIGRP scaling model would work for each one. The advent of Gigabit Ethernet cross-connections to our Telco suppliers, our own multi-gigabit LAN backbone switch capacity, and the performance and stability our Cisco router IOS code delivered, meant that we were able to transit all traffic very efficiently through that central POP. And so, without stress on the network, and plenty of router CPU to spare, the IP addresses EIGRP was asked to handle continued to climb. Which was no problem for a long time. But when the network size reached almost exactly 50,000 networks (mostly 50,000 /32 customer connections), some adherent behavior began to emerge. At odd times we would notice a 5% router CPU increase across all routers, and also the router free memory would decrease. This might last for several hours, and then disappear just as suddenly as it appeared. Examination of the router logs showed the cause was due to EIGRP flapping due to one router becoming stuck in active (SIA). There was however never any clues as to what was causing that, it was always a different router that initiated it, and would hop from router to router until the problem fixed itself. Many hours were spent trying to track down the cause, and I don't think there was a page on the CCO site that referred to EIGRP that was left unread, nor did we fail to bug the TAC with the usual three or four fault cases to try and narrow the field. To cut a long story short, I ended up having to conclude that there was no specific fault. Rather, we had just 'discovered' a previously unknown scaling limitation of EIGRP. And I had to face the nerve racking task of implementing a change of routing protocol, in real time, on our core production network. That was a year ago today. The change plan went to schedule. We were in fact able to complete the task with only a few seconds interruption to most customers. The nightmares I had imagining what could go wrong during such a change now seldom reoccur. But what replaced EIGRP as our IGP? We have a suggestion made by a sales engineer from Redback/Ericsson to thank for that. Since, apparently, their implementation of OSPF was resource hungry, they recommended using BGP. (We were evaluating a Redback router at the time, but that is another story). I was skeptical that BGP, a fine EGP, but I thought unwieldy as an IGP, would work. Initial trials proved me wrong. After a year of faultless operation, it is, I think, the best IGP.

Posted by Steve Waddington at 15:57

Friday, August 31, 2007

Voodoo

"I know it's not working, but it's not Voodoo" I said as I was explaining to a customer, the Financial Controller of a six site VPN network we were configuring, where the cause of his problem may be. The root cause of the problem was something that has long disgusted me and is a major, if not the major, problem with the IT industry in Australia. It is where so called professionals pass themselves off as 'experts' in some particular field, but rely solely on the competence of others, usually the supplier, to deliver a service to their mutual customer. These 'consultants', 'solution providers' or 'systems integrators' as they variously like to call themselves are just a parasitic fact of life and are only a minor annoyance in most cases. Now don't get me wrong, there are plenty of real consultants, solution providers and system integrators that provide their clients, and the industry, with excellent service and value for money. They can be a true partner to a small to mid size company, providing timely expert advice and assistance. My gripe is the ones that ape the true professionals, taking the money and providing nothing of value. As was the case in point. A problem, one of many during this network setup, had been struck and a particular host on a network could not be reached. This had happened before, and in each case our engineers had patiently worked with the on-site consultant to diagnose and fix the problem. Each time it had been due to some local area mis-configuration. Incorrect default gateway settings in one case, wrong netmask on some PC's in another case, devices not even connected to the right LAN in other cases. Normally, that sort of situation is not a problem, and any reasonably competent IT person on site can correct it. In this case however, it seems the on-site consultant of choice lacked both the reasonableness and competency required to make the basic config changes needed. Instead, they chose to report to their client that the supplier (us) was at fault. Now the person in charge of this project on the customer's side was, and is I believe, a reasonable guy, and probably a very good financial controller. Who, as you can imagine, knows very little about computer networks. And in good faith, he believes what his consultant tells him. Unfortunately, the fact was, that from our side, apart from telling the consultant what we had determined the problem to be, there was nothing we could do to fix it - short of visiting each one of the customer's sites ourselves and doing the job the consultant was being paid to do. Not something we could really contemplate doing. The consultant continues to bungle the job, and even though the VPN has been checked and rechecked, and the WAN works perfectly, the customer remains unable to communicate with critical servers. Both our network engineers working on the project, and the customer become increasingly frustrated. The customer, based on the reports from his consultant of what he believes to be non-delivery of service on Exetel's part, positively loathes our Network Engineering Manager. And it doesn't take very long from there before something along the lines of "I demand to speak with a director of your company" reaches my desk. Having had a heads up earlier this could be a problem, and having offered some advice, I was already mostly across the issues and it only takes a few more minutes to get fully up to date. I call the customer and quickly realize I am talking to a reasonable person, with not great technical knowledge, who clearly knows what he wants, but is not in possession of all the facts. I find the best strategy to follow in these situations is to stick to the key points and state the facts as I know them as concisely as possible. Very few people react badly to this, although they may not like what they are hearing, most recognize they are getting a very straight deal, and appreciate that. Hence, part of my explanation was, that even though it might seem like it, computer networks are not voodoo. They follow understandable principles of operation, which allows for accurate determination of where the problem is when something goes wrong. Knowing nothing about any technical aspect of networking, but given conflicting reports from two parties, how is one to tell who is right or wrong? So I offered this advice; that if Party A says "There is a problem, I don't know what it is, but it must be the other guys fault". And Party B says "There is a problem, I have diagnosed the part of the network we control and as far into the rest of the network I can reach, and narrowed the fault to this specific area". Then the least specific set of statements must lose out to the more specific, if no more specific explanation can be offered. The blame game these charlatan consultants like to play is readily exposed to such logic. Their trade is in generalities and vagueness, and their facade can't withstand the light of clear analysis. I guess you can gather, I really don't like those sort of people very much. There is no place for their voodoo in this industry.

Posted by Steve Waddington at 19:48

Thursday, August 30, 2007

The Zen of Network Management

I was asked by one of our junior engineers today, in not as many words 'What is the best way to run a network?'. He was referring to the almost zero downtime and succession successful (ie zero or low impact) changes and additions to the network. We have grown from our first two routers and a single switch in February 2004, to nearly 100 discrete, major, managed devices that have in various ways interdependencies with one or many other devices on the network, and across now six locations and three states. And the network has been grown, component by component, with not more than ten hours of scheduled maintenance and five hours of unscheduled downtime in three and a half years. I can just imagine a new support engineer starting at Exetel, perhaps in their first industry job, looking around for the complex nerve centre, impressive array of monitors and control desk bound engineers that keep the whole shebang running. 'Has to be on another floor' they must think as they scan our single floor open office in North Sydney for the NASA Mission Control Centre they envisaged. I imagine also the vague sense of disappointment they feel as the reality dawns that no such thing actually exists. At Exetel at least. Which sometimes leads to the question 'How _is_ the network managed?', or sometimes the question above, since it is obvious the way it is run is different to the way many people (and many other ISP's I expect) think. Each time I am asked that sort of question, I use it as a checkpoint to consider if what we are doing is in fact (still) the best possible way. Always I will give the questioner my somewhat lengthy paper on Network Management and ask them is they can see a better way of doing it. Often the answer is 'no' and I have a new convert. Sometimes the answer is 'yes' and some improvements are made - which is a fantastic outcome - a great method is made better, and not only is there a new convert, there is a new convert with some ownership. And nothing beats that to encourage dedication in an engineer. The crux of truly great network management (the Zen) comes down to three actions: 1. In every case where there seems to be a need to change the network, do nothing, but consider why it seemed there needed to be a change 2. In the case where there continues to be a need to change the network, do nothing, but consider all the repercussions of such a change for as long as possible 3. In the case where the time frame for consideration has elapsed and there continues to be a change, make the minimum change possible to the least amount of equipment to effect a resolution to the situation. There is an underlying foundation needed to be able to adopt such a policy, and that is the network design must be 'good' in the first place. 'Good' meaning structured, scalable and overall consistent. 'But how can you just do nothing!?' some people exclaim. If you consider any network of reasonable size, almost all 'faults', excluding those of equipment failure, are caused by some well meaning (or not so well meaning, or careless or just incompetent) engineer making some change. In fact, twenty years of measuring these sort of thing has shown me that well over 86% of all network faults are caused by some action or change made in the recent past. A network of reasonable size is also a network of substantial complexity and interdependence. A change at one point may, and often does, have repercussions at another point that may seem completely unrelated. Unless that is fully understood, taking another action to fix the apparent problem will at best have no effect, and at worst cause yet further problems. Even in Exetel's brief history, there have been times in the dozens where some event on a supplier network has cause an apparent fault on the Exetel network. Rushing in and changing code or rebooting routers may seem like the right thing to do, if you are the duty engineer without the experience to know how wrong it is to rush in and 'DO SOMETHING NOW' to fix it. But in each case, it has always proven much better to 'do nothing' other than take the time to consider just what is the root cause, review the change logs and the supplier maintenance notices. I can't think of the last time where doing nothing has not in hindsight proved to be the best first thing to do.

Posted by Steve Waddington at 14:59

Friday, August 24. 2007

Working for a Small ISP

I get on average, one or two job applications per week from people seeking technical roles, and more usually employment as a Network Engineer. And it is not unusual to see a resume from someone with exceptional qualifications and experience. It intrigues me what challenge and progression in their career they believe a small ISP like Exetel can offer them. Small ISP's can in fact, and we do, represent an opportunity to gain practical and hands on experience for someone at the start if their career that would just not be possible in a larger organization. But what about people who already have ten or twelve years experience and a fairly distinguished track record, such as an application I received yesterday? It got me thinking about what technical innovations we have accomplished at Exetel over these three and a half years, and you know, when you list them, they don't look too bad at all: Asymmetric Bandwidth Control - Port based bandwidth controls using minimal router CPU, isolation of traffic using IRP. Integrated Process Network Management - Time of day and periodic routing policy control, automated through server cron jobs. Virtual Training - A method of running a multi-day training course for staff independent of geographical location, achieving measurable results. Interactive FAQ - the structure and management of an AI to answer customer broadband questions. Distributed Network Monitoring and Management - Eliminations of the 'centralized NOC' concept while maintaining minute by minute control of the network. Cost Effective Traffic Redirection for End User Urgent Notifications - Most effective for failed payment, spam and copyright notifications. Modular Scaling Network design and buildup - Scaling a network to any size in low cap-ex impact increments. Those developments have enabled us to save a great deal on operating and operational costs while maintain very high production quality. But then what is the point of innovation if it can't do that? Do other small ISP's do that? I don't know, but they must have some equivalence of efficiencies, or they would all lose money matching market prices in this very competitive industry. So sure, if I was mid career and wanted to work in the best environment to work and innovate at the leading edge of network development, I would find Exetel attractive. But until I listed them here, I don't think anyone outside of Exetel (and even some within) would have known about it, so how did the people who apply for jobs here know that?

Posted by Steve Waddington at 09:42

Thursday, August 23, 2007

Functional vs Flashy

We have commenced a restructure of the management interfaces to our business products today. After three years (of faultless operation mind you) I guess they are looking a little tired. When I first set up those services in 2004, I drew on my experience of running commercial hosted services and thought it would be possible to offer such services with a customer management interface that was 'bare bones functional'. That is, the end user had exactly the tools they needed to self manage their service via a vanilla no frills web portal, and no more. There were two reasons for that. First, frills add complexity and complexity reduces reliability. With only what is needed and no more, I figured one potential cause of any downtime could be eliminated. Second, dammit Jim, I'm a network engineer, not a web designer. I do think however that a web interface can encompass a lot of complexity and still retain elegance - that is, enhanced features but intuitive to use and at no sacrifice to robustness. With the team of very skilled people we have now, it is something within our grasp to achieve. Our plan is to make the portal available listing all the features we want, but not necessarily active. And then progressively activate the features as they are developed. It will allow us to improve what is there quickly, get early feedback from our customers so we can direct our development resources to getting the popular features completed earliest.

Posted by Steve Waddington at 13:53

Tuesday, August 21, 2007

BGP Load Balancing

I am training a new engineer on BGP today. BGP load balancing is critical to our ability to maintain optimum bandwidth levels at a cost that doesn't send us broke. There is a type of reverse logic thinking needed to understand BGP that seems to vex some people, but it's pretty straight forward. Let's call it "The fundamental theorem of BGP" - The networks advertised out determines the traffic that comes in. The networks accepted in determines the traffic that goes out. BGP does not give particularly fine grained control at small bandwidth volumes, but is great at larger levels. The largest network mask accepted in the global routing table is /24, which for practical purposes equates to the bandwidth of 254 end users. There is some variance from /24 block to /24 block, but on average, a fully utilized /24 is going to pull 5Mbps of traffic. So by advertising specific blocks to one supplier or the other, traffic load can be shifted between suppliers in 5Mbps chunks. Which, as our current 1.5Gbps of ingress bandwidth equals .3% fine control. Not bad at all. Of course it would be fairly tedious to shuffle bandwidth in .3% units with a differential between suppliers of 10%. Smaller prefix's is the answer of course. We can currently split the traffic as follows: 16 x /24's - for fine tuning of 5-80mbps 1 x /20 - gives a single 80Mbps block 1 x /19 - a 160Mbps block 1 x /18 - a 360Mbps block 2 x /17 - only for historical reasons - our first APNIC allocation was /171 x /16 - our current APNIC allocation on the 220.233.0.0/16 block. On a day to day basis the bandwidth is fine tuned to within 10Mbps from both suppliers. Hourly variances make it impractical to go any further.

Posted by Steve Waddington at 09:02

Monday, August 20, 2007

The Sky is Falling (Yet Again)

For as long as I have been building ISP networks (since 1994), about once every 18 months there has been some prediction of 'the end is nigh' for the Internet. The latest one is this published yesterday. As always, senior people draw on impressive statistics to show that 'some thing' will cause the end of the Internet as we know it. This time it is IPTV, other times it has been IPv4, DNS root servers, Spam, NAT, etc, etc. Yet the Internet, somehow, endures. Those people should know better by now. Maybe they haven't been around long enough (or have forgotten) how incredibly reliable the Internet actually is. Maybe 'reliable' is the wrong word - as individual networks that comprise the Internet vary greatly in their reliability. Fault tolerant may fit better, or perhaps 'meta reliability'. For this current particular cause of the Internet's imminent demise, it is in fact something that should be of concern locally to any non-global tier 1 ISP. And it has been foreshadowed for a year now by the growth in P2P hand in hand with the dramatic increase in bandwidth capacity to the end user through ADSL2. And, for about the same time, we have been investigating just how we can manage the increased cost in bandwidth and network capacity such traffic generates. The solution is almost certainly going to be some form of caching - and old friend from the 90's I haven't had the need to visit for many years. I will dig up the 1997 Logan-Waddington paper 'Internet Traffic Caching Outside North America' as the principles then will still apply now (unless someone has changed the speed of light). Hopefully the Internet will still be around by the time I find it.

Posted by Steve Waddington at 14:22

Saturday, August 18, 2007

Modular vs Monolithic

Through 12 years of designing ISP networks, the network model I have come to favour is one that of modular-scalability; that is, it starts with a small number of components of relatively small cost, and then more components can be 'plugged in' as the network traffic/connections increase. The Exetel network (a fifth generation refinement of M-S design) for example started with just two Cisco 7206 routers - one was an LNS and one was a border router which took us to 4,000 connected end users and 150Mbps before another router was added. Now after three years there are two 7301 border routers, two 7301 core routers four 7301 LNS routers and two 7206 LNS routers which carry 1.5Gbps of traffic and service 40,000 connected end users. In terms of routing power, we could have achieved the same thing with one or two much larger routers, say a Cisco 10,000 series or one of the Juniper range. However the advantage of M-S is that the initial capital cost is much lower, as too is the incremental cost when capacity has been reached on a particular device. The best way to illustrate this is to consider the price over capacity of a midrange router like a 7301 - when we first started to buy them in 2004 they were \$23,000. As an LNS a 7301 can handle 8,000 DSL connections with some room to spare, or as a border router, can handle 900Mbps at 65% CPU. A monolithic router, such as something in the Cisco 10,000 class, might take 40,000 connections, and in that configuration cost in the order of \$140,000. Whenever I have run those figures, it is always the case that the cost/capacity ratio of a mid range and a large router is almost exactly the same. The difference being that for the first connection, or the first 4,000 connections, it is far more cost effective to spend \$23,000 than \$140,000. Then, for the 4,001st, 10,001st, 20,001st etc connection the impact of another \$23,000 on cash flow is far better for a small business than \$140,000 for the 40,001st connection. There are however some not insignificant intangible advantages of a monolithic model, such as management, footprint and route policy control. As a network continues to grow, the advantages of one begin to outweigh the other. The overlap of where the two are 'about equal' in cost/benefit is quite broad however (thanks to the very similar cost/performance).

Posted by Steve Waddington at 14:59

Wednesday, August 15, 2007

Desktop OS

Today is 12 months since I took a deep breath and changed my desktop from Windows to Linux (Ubuntu). Here are the main plus and minuses I have found using an open source OS for the last year:+++ much better spam filtering for email+ mbox better than Outlook email format++ much improved email retrieval and sorting- available email clients not quite as smooth as Outlook-- winmail.dat attachments not easily readable+ Open Office works just fine in most cases- Open Office spreadsheet is slow for large cell numbers- Some Open Office document formatting is a bit off- printers are more difficult to set up+++ common CLI utilities for networking are much better++ can double as a pre-development development server- Dia used for network diagrams is not compatible with Visio diagrams+ dig+ vim++ grep+ you can understand what is happening with windows share permissions+ TSC works very well and supports more than one protocol++ Synaptic package manager/Apt-get- uninstall-- installation of non-package manager applications more complicated++ Ubuntu how-to's and forum help is helpful- a light bulb in Open Office instead of a paper clipAll up - 22 pluses, 11 minuses. I am 2:1 in favour of a Linux desktop for anyone in a network engineering or similar technical role.

Posted by Steve Waddington at 15:55

Job Goals

We are starting to introduce structured job goal planning this month. I have used this particular method since 1995 when I was first introduced to it by my then boss and now business partner John. It is often pointed out to me that managing people - getting the best from them - is an art rather than a science. Be that as it may, as a literal minded engineer, I have taken a lot of comfort in the past by having a formalized structure on which to base the most difficult, and most critical for success, of all tasks in business. The timing of introducing this method to the business now seems 'about right' to me. Certainly for a company the size we are at the moment. formalized job goals are not strictly needed. But, far better to introduce the discipline now, ready for the time (in the not too distance future) when they will be needed.

Posted by Steve Waddington at 11:03

Tuesday, August 14, 2007

AI

Since Exetel commenced selling broadband three and a half years ago, we have had the idea that an AI engine would be able to answer many of the repetitive questions the helpdesk receives every day. The reason to use an AI engine over a FAQ, Wiki or help section is because of its ability (we hoped) to answer natural language questions, and either guess the correct answer or direct the user to the right question in cases where the user is not sure themselves. Typically, an ISP's FAQ and/or help files grow to many, many pages, creating the paradox that a non-technical user "the people who most need the help - needs technical knowledge to navigate the large volume of information in the FAQ. An AI addresses that paradox because the user phrases the question in the same way they would to a helpdesk person, and, if the AI concept is proven, receives the correct information to help them with their problem. Since the AI knowledge base is written by one or more people, the pattern matching algorithm by default picks up a 'personality', giving the end user the feeling they are talking to a person, or at least, something friendlier than a help file bullet list. As a matter of fact, the end user is talking to another human being, albeit via a time shift created by the AI's proxy of the question response. Building the AI knowledge base is no simple task however. And even more daunting is maintaining it, as there are no standard set of tools to either categorize the knowledge or index the information. For most of the three years the AI has been grown using a combination of 'brute force' (matching every specific question with a specific answer) and trial and error of ways to efficiently manage it. Two months ago I began the latest iteration of optimization of the knowledge base. A large part of that work involved undoing some of the errors, or false paths, we had followed in previous iterations. Early on, for reasons of maintainability, we 'lobotomized' the AI by removing most of the personality files so that they did not interfere with its main task of accurately answering ADSL questions. The good news is the AI now has 62,148 unique knowledge base entries and post the last update, matches better than 77% of broadband support questions. Near future upgrades will include: a redesigned interface, flash animation, re-introduction of 'personality' traits, real time responses to one-on-one support chat via IM. The results of the latest update are encouraging. The ultimate goal is to make the AI indistinguishable from a human response to the 85th percentile of most asked questions. I think we are about another year from that goal as it stands now.

Posted by Steve Waddington at 15:03

Friday, August 10, 2007

Network Engineers - The Good, the Bad and the Ugly

I was having a discussion with our Network Engineering Manager this morning and one of the points was I had noticed an email from a customer who said he was expecting a call (from one of our engineers) on Friday, but hadn't received one. When I see an email like that, I always like to look into the reason that caused it to be sent - it can at worst be an early warning for some bigger problem (which can then be addressed), or, as in most cases, have a once off cause that is unlikely to be repeated. But the main point is, the discussion led to 'what makes a good network engineer?'. As I pointed out, anyone can read a manual, enter IOS in a CLI or complete some industry recognized exam. Which is great.

However the ability to do those things makes them no more a good NE than a teenager is a good driver just because they have passed their driving test. There is one thing that truly differentiates a good, in fact great, network engineer, or anyone in any technical role for that matter. The one thing that the lack of can make dealing with technical organizations a total nightmare. The one thing that I am very sad to say is totally lacking in some of the companies we deal with, who it would surprise no one to reveal are large Telco's who should know better. "Ahad" I said "errors can be forgiven and bugs can be fixed. But what makes a really great engineer really great is that they always keep in mind where their salary comes from". What I meant by that, as Ahad, who is a paragon of that virtue, understood, is that every network engineers salary is paid for by the customer and the customer alone. Unhappy customers = fewer customers = less revenue = no need for engineers who create unhappy customers. There are of course many things that will make the customer unhappy that a network engineer has no control over. Neither are technical people particularly good at empathy, nor should they be necessarily. It is often a case of either being good at talking to carbon or good at talking to silicon, but not both. Making a customer happy to whatever extent an engineer is able, can, in my experience, be achieved by simply never making a commitment to a customer that can't be kept, and then making sure that the commitment is kept. If a phone call is promised at 4pm, the phone call is placed at 4pm. If a status update is agreed by noon Thursday, then by 11:45am Thursday the update is on its way. All the best engineers I have known have had or acquired that discipline.

Posted by Steve Waddington at 15:51

Monday, June 4, 2007

Wifi in the US

Since 1998, a handy feature of each Interop event was the availability of free, and good, wifi access to the Internet. Anywhere you were at Interop, you could just flip open your laptop, check your email, access your network or whatever you needed to do, fast and reliably. So too, since at least 2001, wifi has become increasingly common at hotels in the US - sometimes for a fee of \$10 per day, or just as often for free. You may also have heard about wifi at Starbucks and Mac Donalds, nationally and freely available for patrons. If that was the case in 2001, then by 2007 it could only have gotten better, and more common, right? Well that is what I thought too. I was right in one respect, wifi is now very, very common. From just about any location there are four or five, or more, public wifi networks, and usually the same number of private networks as well. Unfortunately, the very popularity of wifi is its downfall. There were so many networks within reach of each other that the spectrum was saturated. Signal quality for any one network might be 5/5, but because of contention with other networks, it was hardly possible to reply to a single email without the connection dropping out for a minute at a time. Very frustrating when using webmail or IMAP to access a remote email server. The last thing I expected was to have unreliable Internet access while traveling in the US, but that is definitely the case today. Dial-up is still very popular for that reason. The only reliable Internet access was at Hotels that had an ethernet cable connection in the room.

Posted by Steve Waddington at 12:15

MPLS

Stands for Multi Protocol Label Switching. Basically it is a network routing protocol that a carrier deploys on their network so that a customer who connects also using MPLS can use part of the carrier's network as if it were their own - that is the customer has control of the routing, VPN setup and routed traffic as if the customer had bought dedicated circuits between each site. The technology isn't that new. Cisco and Juniper have been pushing it as the way carrier networks must go since the late 90's. I investigated it at One.Tel and came to the conclusion it was a technology the market was not yet ready for (in 1998). The Swiftel network was all MPLS - but still in 2002/3 no one (customer who might use it) understood, or cared, what it did. This year at Interop, both Sprint and Verizon proudly announced MPLS as the latest and greatest thing customers who wanted VPNs should buy. So of course, I was expecting to have to elbow my way through a crowd 20 deep to get to the Sprint stand to find out just how they had achieved such a breakthrough. For some reason though, I found myself almost alone at the stand with two or three other people, and five Sprint reps. Obviously the free Sprint t-shirt give aways were not in fashion this year. The Sprint rep I cornered, once he heard my Australian accent, wasn't that interested in talking to me since it was clear he wasn't going to sell me anything, but called over one of their 'solution consultant' engineers when I asked what other protocols besides IP they would support. Mike, the solution consultant, had just been on the Sprint product course run by Cisco, and so was brimming with information about the 'MPLS revolution' about to sweep the world by storm. I must say, in truth, my expectations were not high about learning of some breakthrough in this technology. But Mike was so enthusiastic, that for a moment I began to think 'hey, maybe they do have something'. Sadly though, I can only report that my expectations were not exceeded. The conversation went something like this: me: Apart from IP, what other protocols is your MPLS network able to carry? Mike: Uh, well only IP really. me: Sure, I understand, does anyone actually want anything other than IP? Mike: Not that I have seen. me: Me either, maybe IPLS or SPLS would be a better name for it, ha ha. Mike: Ha ha, yes I guess. me: So with this MPLS, what routing protocols would the customer use to control their traffic? Mike: BGP. me: Only BGP? What about an IGP like OSPF or even RIP? I hear Cisco now offers EIGRP support as well? Mike: They do? I mean, yes, they do. But no, we only support BGP. me: Oh, ok, so how does the customer figure out how to use BGP? Mike: Well most wouldn't know, but we provide that as a managed service. me: That makes sense. But if you manage the service, doesn't that remove a lot of the flexibility of MPLS in the first place? Mike: Well yes, but the customer tells us what they want at set up time and we build the circuits for them. If they want to change at any time, we can do that but there is a fee for service. me: Of course. How often do you find that customers want to change? Mike: Not very often. Usually once it is set up the only changes are adding new locations, but the basic network design doesn't change. me: Yes, I find that too. me: But of course, the other great flexibility of MPLS is the dynamic bandwidth in a fully meshed network available to the customer, only limited by the local loop capacity they connect with. I would think here in the US where bandwidth is so cheap, that must be a very nice feature for a customer to have? Mike: Well

yeah, but we don't let customers just use any amount of bandwidth you know. They say what they want for each site and then we limit the circuits to that.

me: Really? So say a customer has five sites, they nominate the bandwidth for each site?

Mike: Yes

me: But then the bandwidth between sites is still dynamic, right? So if site A wanted most bandwidth to site B most of the time, but then one day need to change to site C, it wouldn't make any difference?

Mike: No, the customer has to say what bandwidth they want between sites if they want it fully meshed. But most of the time they just want bandwidth back to a central location so the network we build for them is 'hub and spoke'. Of course if they want to change that we can do it for them, for a fee for service.

me: I see.

me: Ok, one last question. Say you had a customer who knew all about BGP and was completely able to configure their routers to connect to your MPLS network. Would they be able to get those benefits?

Mike: Not right now, we plan to release that some time later one, but for now we only offer it as a fully managed service. We don't think the market is ready for anything else right now.

me: Look Mike, thank you for your time explaining that to me. I don't want to be rude, but isn't what you described just a standard VPN hub and spoke network, the same as if is were built with GRE tunnels?

Mike: Oh no, not a all.

me: How is it different then?

Mike: Well this is MPLS!

me: I see, so your network goes to eleven then. Thanks for your time.

Mike: No problem.

Americans aren't big on irony, so I think the meaning of my last comment was missed.

Posted by Steve Waddington at 12:12

Sunday, June 3, 2007

Hot Products from Interop

The 'hot button' products in IT are: enablement of better collaboration within the organization, enablement of better collaboration with external suppliers and customers, doing that securely, reducing the cost of communications but not at the expense of reliability. Cisco has arguably the most coherent bundled solution to meet those wants, through the integration of every part of the network from the ASIC chips to the application layer. However, the marketing spin is one thing - it is still hard to identify exactly what the product(s) is/are that will do that, other than an open cheque book to Cisco. It is important to define just what 'collaboration' and 'enablement' mean, beyond the marketing term, or the picture it raises in peoples' minds as to what it may mean. Collaboration in this context is 'rich' communication with exactly the person or people you want, at exactly the time you want. Rich communication is triple play (Voice, video, Internet) or quad play (Voice, video, Internet, wireless. But not called four play for some reason). Enablement are the technologies or product bundling solutions that allow that to happen. IT departments are being told that they need to embrace web 2.0, they need triple (or quad) play enabled on their network and they need to embark on the next wave of Internet integration of their applications and hardware base to do that. But what does it all mean? John Chambers, President and CEO of Cisco (who incidentally have recently patented 'Triple Play') seemed to know exactly what it meant. But I suspect the waters become a little murkier past him. What can be said however, is that whatever 'collaboration enablement' is, the deliver mechanism of any of those services is an IP network and the Internet.

Posted by Steve Waddington at 12:06

Saturday, June 2, 2007

US Internet Speeds

I asked a few of the people I spoke to at Interop about their typical home Internet speed. In the US for the 'tech savvy' the most common form of Internet access is on cable as part of a bundled deal with the cable company. DSL is nowhere near as common as it is here. Some people had fixed IP addresses, which cost them an extra \$5 to \$20 per month, most did not. I couldn't get a real view of the cost of Internet, because of the bundling, but it seemed to be about US\$25-30 per month for a basic package with 'unlimited' downloads. The wholesale price for bandwidth in the US for small volumes of 10Mbps is \$10 per Mbps - probably why unlimited downloads are no problem. The speed reported by users of cable was highest of 9Mbps, lowest of 500kbps for normal traffic - the average from my quick poll was about 1Mbps. EVERYONE said that p2p speeds were around 100kbps at best. No one was particularly upset about it. It seemed common knowledge that ISP's in the US restrict p2p traffic severely, though no one advertises it. Like queuing for an hour for anything and flight overbooking, it was accepted as just the way things are.

Posted by Steve Waddington at 12:00

Friday, June 1, 2007

Interop Las Vegas 2007

I try to go to Interop or some US based conference once every three years or so. In this case it was four years since I was last there. The strong Australian dollar made it particularly attractive this year, and also I thought one of the courses - CIO Bootcamp, would be of high value. Interop is run twice a year in the US, and once a year in other countries like Japan, Russia and Germany. The US shows are the premier event conferences for the Data Communications industry - much more so than CeBit, which tends to be more general IT in nature. In particular, Spring Interop in Las Vegas is the biggest conference of the year. The reason it is held in Las Vegas is because at one time, Las Vegas was the only place in the US that had a convention centre big enough to hold all the exhibitors. Interop has about halved in size in terms of exhibitors since the last time I was there in 2003. I think that is partly because companies like Cisco keep acquiring smaller companies that would otherwise be there. Never the less, the floor space taken up by the exhibitors there is still about three times the size of the Sydney Exhibition building. Many companies announce new products at Spring Interop, specifically holding them back for the event. The conference is typically attended by VP's and CIO's of the largest technology user companies in the world, such as Wal Mart, Citi Bank, every Telco you could name, NASA, Dept of Defense, and many more. They all go because the event has the latest products, announcements and the best speakers. The speakers and exhibitors (450 this year) go because the event attracts the top decision makers from the largest organizations. I go for all those reasons plus one other, which is to find out where my understanding of technology stands in light of what everyone else is doing. I don't go expecting to find some new revelation about how networking should be done (though more often than not one or two good new ideas come out of the event, but that is an unlooked for bonus). Rather, I gain peace of mind that we are on the right path, on the edge of the curve (if not ahead of it), and are not going to be caught flat footed by some major change in the industry that would leave us struggling to remain competitive.

Posted by Steve Waddington at 11:15