

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 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
- * and 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