The ISP Column
A monthly column on things Internet

July 2010

Geoff Huston

Wired vs Wireless

I never thought I'd see the day when the difference in capability between a wireless and a wireline Internet would become a core policy differentiator in a national election, but this has now happened in Australia.

Perhaps it's a timely indicator of just how important the Internet is in our daily lives these days, and how much we've managed to associate keeping in touch with family and friends with tools such as Jabber and Skype, and just how much of our daily working life is now mediated by email. It seems that everyone has an interest in a ubiquitous, fast and cheap internet. Now that interest has been taken up as a major policy differentiator by both sides of the political spectrum in the recent Australian election. What was this all about?

On the one hand there is the National Broadband Network (NBN), an ambitious project to replace the now quite old telephone copper pair access network with a comprehensive fibre optic system. The copper pair network was originally funded by the public purse many decades ago, and has since been passed into private hands, along with all the other network assets of the former monopoly telephone operator. The NBN plan is to provide a layer 2 national fibre access network that would provide a last mile conduit to the majority of the nation's 6 million households and business premises. The NBN was intended to be revolutionary in terms of the change in capability of the national network and lift achievable last mile access speeds from DSL-speeds of 1 to 10 Mbps to a uniform level of access speeds of 100Mbps for every customer. Curiously, the election campaign has managed to squeeze more capacity out of the network and the political rhetoric has managed to up the access speed of this network tenfold to claimed access speeds of 1Gbps. This network is intended to be truly prodigious and the harbinger of bountiful benefits for everyone for many decades to come! Of course such a massive undertaking to rewire an entire continent does not come cheap, and the budgeted cost of providing this infrastructure to its 21.5 million population is a $43 billion impost, or a cost of $2,000 for every Australian resident. This project is a public works program, and the evolving expectation is that the network will once more be a public asset, in the same way that the original copper telephone network was constructed using a funds underwritten by the public purse. In every respect the project is intended to be game-changing. The ubiquitous use of high capacity across the entire population is intended to alter the way in which services are delivered, in which we define work and entertainment and the way in which a relatively small population in the south Pacific Ocean defines its place as a developed and hopefully highly competitive economy in a global context. These are indeed great expectations and the price tag is entirely commensurate with the level of euphoric optimism that is associated with this national project.

On the other side of the political spectrum there is a proposal to scrap this scheme immediately after the election. Lest this political party be perceived as technical troglodytes, they propose to replace it with a program of installing a swathe of wireless access points, particularly in the semi-rural areas of the Australian continent, and undertake some form of unspecified upgrade of parts of the existing copper pair network. This is a far more modest program, which is reflected in the price tag, currently estimated to cost the Australian taxpayer a mere $6 billion Australian dollars.

Each side of this political debate is keen to paint their chosen Internet technology in the most positive of lights, and portray the alternative in as dark a light as possible. A national network

based on expansion of wireless infrastructure is portrayed as retrograde and hopelessly ineffectual in terms of national infrastructure. A fibre optic network is portrayed as being wastefully extravagant, unnecessary, and behind the times in today's i* world of wireless access devices. As a result, this ordinarily somewhat dry debate conducted between engineering, product and business line managers within the industry about the relative merits and weaknesses of mass access wireless and wireline networks has come to the surface of the political world for a day or two of mass media focus.

What lies behind this debate? In our efforts to convert every home and office into a wifi hotspot and convert every handset into a 3G wireless client have we really turned our back on the wired Internet? Is the copper pair, and even the concept of the fibre access network already being consigned to the dustbin of historical technologies, and will wireless totally dominate the future of the Internet? Or does the wired network have an assured future as an essential path to higher capacity and greater diversity and utility of the Internet, while wireless is just a passing fad that cannot sustain the full load of the diversity of needs of tomorrow's Internet?

The case for comprehensive uptake of wireless networking in the world of consumer electronics is close to overwhelming in today's environment. This month it has been reported that the 5 billionth device will "plug" into the Internet (http://www.networkworld.com/news/2010/081610-5billion-devices-internet.html). It is statistically likely that this 5 billionth device will not exactly "plug in" to the network but wirelessly associate with a nearby base station! Wireless has been focal point for the Internet's evolution in the past few years, fuelled largely by the market success of Apple's various i-devices and the competitive responses from other suppliers. An earlier press story, again from Australia, illustrates this rather dramatic growth of the wireless market sector

> Use of wireless broadband services mushroomed during the past year [2009] to reach more than 2 million subscribers, driven by the popularity of wireless modems and mobile devices such as the iPhone. The Australian Communications and Media Authority's communications report [for 2009] revealed the use of wireless broadband services jumped by 162 per cent in 2008-2009. ... Wireless broadband subscribers accounted for 25 percent of the number of Internet subscribers, up from 11 per cent in 2008."
>
> *The Australian, Wednesday 13 January 2010*

There is no doubt that wireless services are so popular with consumers that they attract a price premium for their services. It is still the case that SMS messages in the mobile network are the most expensive data service on the planet, measured in units of dollars per megabyte, but other wireless 3G data services are also up there in terms of the margins of price over cost, particularly if one is daring enough to use international mobile roaming services for data! In the Australian market, for example, for the same $50 per month a consumer can access an Internet service with a usage cap of 3Gb per month with a wireless service provider, or take a service with a cap of up to 100Gb per month with a DSL service provider. Why is the wireless service some 30 times more expensive? The cost or provisioning a wireless service is dramatically lower than the cost of a wired service. The return on the investment of a wireless tower in densely populated urban environments is dramatically higher than the business case of dragging more wires through existing communications conduits. Even taking into account the lease costs of the radio spectrum, wireless services still represent a much higher margin activity than wired services. So its evidently not the relative costs of the service that determines the retail price of the service. Perhaps its more of a case of provider push coupled with consumer pull. Wireless services resonate with consumers in terms of convenience, and are prepared to pay a premium for this convenience. Consumers are prepared to pay higher prices for mobile services. Providers use this preference to add a premium to their mobile products and services, making this a more attractive product for them in terms of return on investment in service infrastructure. In every respect this looks like a mutually satisfactory setup.

There is always a "but" in these arrangements. The perennial question that gets posed about these innovations in service delivery in the internet is: "But does it scale?" When we consider not

just a population of some 6.8 billion humans, but a population of over 100 billion chattering devices, will this approach scale?

In other words, is wireless an infinitely exploitable resource? Can we expect ever-increasing numbers of services, ever-increasing intensity of use, and ever-increasing capacity from the wireless network in the future? Like the air we breathe, the radio spectrum is a shared resource with many competing potential uses, from broadcast media, such as radio, television and geolocation to private point-to-point services with mobile telephony, and various permutations of satellite services. And of course, not all radio spectrum is the same. Lower frequencies provide better penetration through buildings, and can extend beyond line of sight, but have limited bandwidth. Higher frequencies have higher bandwidth, but are more readily absorbed by hills, buildings, and even walls. And of course the radio spectrum is a shared medium, so it is necessary to manage the spectrum as a common resource and coordinate the various potential users of the spectrum. The spectrum space is already full, and the prospect of displacing the existing users to make way for a massively larger Internet appears to be an unlikely outcome (See http://www.newscientist.com/blogs/shortsharpscience/spectrum.png for one view of the existing spectrum allocation).

The part of the spectrum is that can be used for wideband digital communications is very limited. As more subscribers crowd in the same shared spectrum space the problem is that the quality of the service ultimately degrades. This can be mitigated to some extent by using more and more base stations with smaller radii of coverage, but at the same time this approach increases the issues with cross talk and signal interference and the complexities with mobile station handover.

The service performance with wireless also suffers, with signal dropouts, higher bit error rates, higher jitter and sudden changes in access capacity due to the method of channel contention in 3G. All of these are particularly hostile to the TCP protocol, and while there has been much said about the rapid rise of theoretical carriage capacity of wireless systems in recent times, achievable sustained data transfer rates using TCP in the wild fall far short of the hype.

Does an investment of $6 billion of public funds into wireless infrastructure represent a wise investment in national infrastructure for a future extending for many years into the future? Or would this be a case of making a investment in a current technology that is closer to a fad than an enduring element of a digital infrastructure? Also, given that the current wireless internet has already been enthusiastically constructed with private capital investment, should further public funds be expended in undertaking a public works program that may well be undertaken by private capital in any case?

There is no doubt that if we are facing a bandwidth hungry Internet future, then fibre optic wireline services can provide much greater reticulated capacity to the network. Unlike wireless, wireline systems behave consistently in terms of bit error rates, latency management and jitter. As a result the TCP transport protocol behaves with close to maximal efficiency, and can achieve sustained data rates equal to the line bandwidth, even at gigabit per second rates. All this prodigious performance can be achieved on fibre optic systems without crosstalk and without interference between users. Because the signal is guided by the wire the systems exhibit far higher energy efficiency, and can operate at far higher speeds. If speed and capacity are what we are after, then speed and capacity is precisely what fibre optic systems can deliver.

But of course despite all these efficiency and performance differentials, it's still a wired service, and the service is tethered to the end of the wire. That limits its usefulness and utility in an acknowledged highly mobile world. However, the choice is not quite so stark as a choice between an RJ45 connector and a 3G modem. WiFi has also revolutionized the consumer marketplace, and these days its quite commonplace to see appliances use WiFi as a connection medium. There is no doubt that I have no interest in using a carrier's 3G network to send a print job to the printer sitting beside me on my desk – that's a job for my local WiFi network, as is access to a home server and a myriad of other local operations in the home and at the office.

Where should public funds be spent? On a comprehensive revamp of the wired access network, replacing the aged copper pair telephone network with a highly capable fibre optic network? Or on improving access in those areas where the copper pair network simply cannot support high speed access by public investment in wireless infrastructure?

In trying to answer this question, we return to a persistent theme in the area of public communications infrastructure. What's the role of public capital investment and how is that balanced against the role of private capital investment? Is it possible for private investment to fulfill the entirety of a public agenda? Given that a capable, cost efficient and effective public communications infrastructure that encompasses an entire national constituency is seen as a core deliverable of any national communications policy regime, then how is this best achieved today?

To move back from generalities to the specifics of this broadband investment choice, is it realistic to expect that we have further decades of useful life from an already ageing copper pair infrastructure? As a consequence, should current public investment focus on current gaps in the national infrastructure, using a relatively cost effective approach of plugging these gaps using wireless infrastructure where the copper network is simply inadequate, and leave the remainder of the network in situ, as being adequate for the moment? Or there is the option to leave such wireless infrastructure investment to private enterprise, given that this technology is enjoying strong consumer attention and there is a continuing investment in wireless infrastructure by the industry actors. Instead, should a public investment program focus on a longer term national program of replacing the copper loop with a comprehensive fibre optic network? From such a longer term perspective perhaps the NBN is the better approach, as we need to concede that the level of investment required for a national very high speed access infrastructure in a fibre access network is probably well beyond the scope of private capital works investment. So far all that the industry has achieved in this space has been the rewiring of the CBDs in the major cities, while the upgrading of remainder of the network has been effectively ignored. It appears that this is, like many major infrastructure projects in the past, one that properly sits in the realm of a public investment program, in the same way that we've made investments in national road, rail and shipping infrastructure in the past.

That's the spectrum of choice between wireless and wired infrastructure programs for a better, faster and broader broadband Internet by the two Australian political parties. The wired vs wireless debate has become a matter for the electorate to decide.

Postscript

Since writing this article, the Australian national election was held of the 21st August. It would have been nice to have cited a clear result of the election and note that when this question was passed to the electorate, there was an obvious choice made between wireless and wireline infrastructure investment models for tomorrows national network infrastructure. But that was not to be, and the result of the election is an evenly poised outcome with no clear decision one way or the other.

It seems that in communications what we really want is both wireless and wireline.

What do we want? Everything!

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net