

April 2013
Geoff Huston

What the Ftt?

Three years ago, in September 2010, Australia held a federal election. At the time I thought that I'd never 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 that was what happened in Australia in 2010. The debate was about whether the country would be best served by a National Broadband Network (NBN), an ambitious \$43B project to replace the ageing telephone copper pair access network with a comprehensive fibre optic data reticulation network that was designed to deliver some 100Mbps of data to each termination point, or 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, at a total cost that was estimated at the time to be some \$6B. The outcome of that election was that the political party that was the proponent of the NBN formed a minority government, and the support of a number of independents for this government appeared to be based, in part, on some broad level of support for the NBN in rural Australia. You could say that wired network infrastructure won the day over wireless.

Now, almost three years later, we seem to be having a national déjà vu moment. Yes, there is another election coming in September, and yes, once more the country's national digital network is firmly in the sights of the politicians and it's a part of the political debate.

The NBN being extensively promoted by one side of the political divide. The costs are slightly higher than the original estimates, which is not unusual nor unexpected for a public project of this dimension, and the timescale of deployment has also slipped somewhat, and again that's not unexpected for government work! The technical profile of the network is largely unchanged. It's a fibre reticulation network that is claimed of being capable of delivering a relatively uniform level of 100Mbps service in a consumer market context to the country's 6 million households, and a capability of delivering up to a gigabit per second through the fibre tail circuits for a relatively small subset of terminations in the business and public sector. The network is a Fibre-to-the-Home (FttH) model, where the grand plan for the aging in situ copper plant is to be pulled off the poles and hauled out of the ground, and sold off by the ton, to be replaced by shiny new fibre optic cable.

On the other side of the political divide their plan has been extensively revised from that of 2010. Wireless for the rural fringes and copper for the cities is no longer a core part of the plan. The new plan is Fibre-to-the-Node (FttN). In this plan the fibre is run out from the exchange to powered "node" units on the curb, with each node serving a few hundred houses. The conduit between the node to the house would be the recycled copper pair. To quote a current news article on this: "The Coalition policy will cost \$29.5 billion, with a goal of giving all Australians access to internet speeds of at least 25 megabits per second by the end of 2016." [<http://www.abc.net.au/news/2013-04-09/conroy-hits-out-at-coalition-policy/4618232>]

So is this really happening again? Are we going to have an election where one of the major policy elements being put to the voting population is a public communications infrastructure plan? (and yes, Australia is one of the few countries where voting in elections is compulsory, so we all get to have our say, like it or not!) It's going to be very interesting to watch the debate, as this time around the critical distinction here is nowhere near as obvious as wired vs wireless, but a far more subtle distinction of whether its to be FttH or FttN!

What can we make of this distinction?

Let's look first at FttN. In Australia the old monopoly telco has certainly been here before. Some thirty years ago the price of copper was soaring. And Australian cities were growing. The urban dream in Australia continues to be a detached house on its own block of land, so the urban density is low and the copper runs from individual residences to the exchange can be quite long and quite sparse. Back at that time, as a cost savings measure no doubt, the national telco came up with a hybrid system of Remote Integrated Multiplexors, or RIMs, for new housing estates. Each RIM serves around 300 houses. (A RIM has some 480 voice circuits in aggregate capacity as a general rule, and the provisioning rule was to plan for 1.6 copper pair services per household, hence a RIM was good for around 300 houses.) The RIM connects to the exchange over a fibre cable, and uses shorter copper pair runs from the RIM to each residence. If all you were doing was voice telephony, then RIMs served telephone customers effectively, as well they should. It reduced the provisioning time for telephone services, and cut the initial capital expenditure of installing long copper runs. So for telephony RIMs were an ideal approach. But data is an entirely different story. At the time RIMs were deployed, in the 1980's, data was relatively esoteric and DSL simply did not exist. If you wanted to access a bulletin board then you used a modem. 14.4Kbps was considered a very fast speed and 28.8kbs out of an analogue modem was unheard of speed! As long as a RIM could deliver at this speeds, then the RIM was not a problem. But as the 1980's progressed and we entered the 1990's our requirement for data continued to rise. By then the 28.8Kbps modems were unacceptably slow, and even the best 56Kbps modems paled into obsolescence when compared to the new wave of DSL deployment into the residential market. But if you were connected to the exchange via a RIM then this newfangled DSL service was not for you. If you were lucky you were connected though an Integrated RIM that offered a 56Kbps digital service, as the RIM system passed the digital signal into the voice exchange as a standard 56Kbps digital circuit as used for digitised voice telephony. However that was a best case. If you were connected to a RIM that used a breakout on a junction frame into individual jumper pairs for each of the 480 individual connections at the exchange end of the RIM service then, by virtue of the RIM design, the best you can do was an analogue modem running at around 28.8kbps - 31.2kbps. In the days of close to ubiquitous use of DSL in the market, consumers who found themselves behind these RIMs also found themselves locked in at glacially slow modem speeds with no clear path to get to the megabit speeds that everyone else was using on the copper pair networks. The newest and most recently constructed housing estates were rapidly condemned into being labelled the worst of the digital slums of the information society. So the first iteration of the FttN approach has not been a pretty story.

Now if the last experience of FttN was evidently so disastrous, why would anyone want to contemplate it again? Part of the reason why RIMs were used in the first place was not because they gave superior telephone service, but because they were simply cheaper than the alternative of long and expensive copper runs (in the ten years from 1970 to 1980 the commodity price of copper had doubled in real terms). And a large part of the reason why FttN is being contemplated again for this country is much the same reason. It's cheaper than the FttH alternative. But in this case what is cheaper is the cost of pulling, hanging and burying cable from the common node termination point to each residence. And Australia has some 6 million residences. If we reuse the existing last mile copper network, then we avoid this additional pulling, hanging and burying effort for some 6 million terminations. And that will save large sums of money. In the FttN framework the consumer uses the existing copper pair network to get to the node, and the digital stream is then passed into a trunk fibre system as multiplexed packets or cells.

So the critical question in the FttN approach is: How do you get from your house to the node over this copper? Well good old DSL, that's how! How fast does DSL go? Like the question "how fast does my mobile data service go?" the answer about DLS speed is "... it depends". In Australia we make extensive use of ADSL2+, a technology that uses 2.2Mhz of spectrum on the copper pair to achieve a downstream rate to the customer of as much as 24Mbps and with Annex M, as much as 3Mbps upstream. So if the package is FttN, and a copper pair with ADSL2+ Annex M between the node and

the house, then what you get at the house is a service that operates as a speed of up to 24Mbps downstream and up to 3Mbps upstream.

Maybe we should ask to other question: What do you get from the existing copper pair network? Well, it's likely in the urban areas and semi-rural areas its ADSL2+ Annex M. Yes, that operates as a speed of up to 24Mbps downstream and up to 3Mbps upstream. So what does the country get for its investment of \$30B into a national FttN network? We would've subtly change the set of probabilities in the "up to" phrase about DSL speeds, and increased the odds that you are more likely to get a stable speed of 24/3 Mbps on your Internet connection. It's hard to call this a clear bargain!

But perhaps I'm being a little harsh on the FttN approach. If we take a step back and look at the developments in transmission of digital signals over copper over the last 50 years, then we have achieved a series of prodigious feats. Using a combination of constant improvements in digital signal processing, and refinements to the signal encoding technique we've managed to get unshielded twisted copper pairs in bundles to pass megabits of data. Today its possible to send and receive 1 Gbps over Cat 5 copper pair cable plant. So why shouldn't we expect that FttN will promise higher speeds in the future over the same copper pair plant? The problems of using copper pairs at ever higher speeds come in two forms: noise and crosstalk. Yes 1Gbps over 4 copper pairs of cat 5 twisted pair is possible, but the longest possible cable run supported by the 1000BASE-T standard is 100m. The longer the copper run, the higher the level of noise. One way to compensate for noise is to increase the power of the signal. But this is where crosstalk becomes a problem. Each copper pair acts like a radio transmitter, and a certain amount of power that was pumped into the copper pair is radiated outward from the unshielded copper pair as radio waves. Copper pairs also act as radio receivers, so if a copper pair is placed adjacent to a pair that is carrying a high power signal, then the pair will pick up part of that signal. Crosstalk occurs when a set of copper pairs are bundled together, as is the case in the telephone network's copper plant. As the wire plant moves towards the node point the individual pair runs are bundled together into larger bundles. And then the issue of crosstalk becomes a significant impediment to increased speed over the copper pair. If we want to go faster than the 24Mbps of ADSL2+, then we could move to VDSL. Instead of a 2.2Mhz spectrum, VDSL uses a 30Mhz signal spectrum for the downstream signal. But this 30Mhz spectrum presents us with a problem, that the wider spectrum means that we have to increase the signal power and decrease the cable length. So VDSL is only defined over the relatively shorter span of at most 1.8km of relatively high quality copper, and even then the power budget implies relatively high levels of crosstalk across cable bundles. One report [<http://www.ospmag.com/issue/article/vectored-dsl-rescue>] claims that "in practical deployments, where multiple pairs are sharing the same cable and are thus causing crosstalk into each other, data rates over a single-pair are limited to no higher than around 70 Mbps for downstream and no higher than around 40 Mbps for upstream, even with loops as short as 500ft [150m]." So this is not looking all that good. Along comes "Vectored DSL", which uses signal processing at the exchange end to pre-load the signal with the inverse of the anticipated crosstalk (just think of rather clever predictive noise cancelling headphones!) The claim from the proponents of Vectored DSL is that this offers short copper pair networks, as would be anticipated in FttN environments, the potential to achieve speeds of up to 100Mbps in combination with VDSL. Now I have to admit that there is a fair amount of unbounded and potentially unfounded technical optimism in touting this particular form of technology. The pre-loading of an inverse crosstalk signal is not the same as the precise loading of the inverse of the actual crosstalk signal, so if the crosstalk noise deviates from some pre-set bounds then its efficacy drops off quite dramatically. It also appears that you can't apply this Vectored DSL wizardry to every pair in the bundle, and some lines would need to be given "priority" over others in order to optimise the crosstalk compensation function. And of course its highly distance sensitive, where every additional meter of copper pair from the node to your house will probably drop the maximal achievable data rate by a few megabits per second. So if there is to be a FttN node in your neighbourhood, then having it located directly outside your house might be the best possible result for you if you want the highest possible capacity and speed from the system!

What about the NBN and its FttH approach? Compared to copper, fibre optic cable is impressive. Ever since Corning's introduction of low loss optical cable into the communications market in the

early 1970's, optical cable has revolutionised this industry. Fibre cable is transparent to light in three bands, at wavelengths of 0.85, 1.3 and 1.5 microns. Each transparent band is around 25Thz wide, and using an encoding system capable of delivering 2 bits per Hz, that implies a total theoretical capacity of some 100Thz from a single strand of fibre! However such prodigious data speeds are difficult to achieve. Indeed creating digital signalling systems that can operate at such speeds is still on the list of challenges for the silicon industry. So how do we use fibre systems for extreme speed? The approach we use is surprisingly simple: we take a number of digital streams and encode each stream into a unique color (wavelength), and then combine each light stream and send the aggregate into the cable. At the other end we use a Bragg Diffraction Grating to perform an analogous function to that of a prism: to split out each of the "colors" into separate digital streams. This form of optical sharing is termed "wave division multiplexing" (WDM) and is used in the trunk networks of all major communications systems. However this form of packing data into a fibre cable is expensive, and in a large scale FttH deployment its impractical to use WDM to share a cable across multiple users.

One approach to a fibre network is to use a dedicated fibre pair to connect each customer to the exchange. This was the traditional approach used by telephone network operators for their larger business class customers. In Australia in the late 1980's the provisioning rule used by the telco was that any data service that operated at a speed of 256Kbps or higher was provisioned using a dedicated fibre pair, and equipment was placed on the line at either end to provision the speed that was actually ordered.

This always strikes me as such a strange business model. It's like building a slow car by taking a specialised high speed racing car and bolting on additional leaden weights in order to slow it down, all to charge the customer a lower price than the price of the original high speed car!

However, when looking at a large scale deployment, dedicated fibre pairs for each service point is also way too expensive. Even dedicated wavelengths using low efficiency coarse wavelength division multiplexing is too expensive. These days FttH deployments use either Active or Passive models.

An Active Optical Network (AON) model is much the same as a conventional routed network. A common fibre path is constructed from the exchange to a node point. The node point contains a digital switch, feeding dedicated fibre runs from the node cabinet to each optical network terminal. AONs are in effect switched Layer2 VLAN networks, where each ONT is provided with a dedicated virtual circuit between the ONT and the exchange. The node switching point is an active network element and may be provisioned with buffers and various policy controls to implement the desired service elements.

The other form of FttH optical networking is a Passive Optical Network (PON). In this case the single trunk signal is optically split and fed to up to 128 ONTs. As all ONTs receive all downstream data, PONs use per-endpoint encryption on the downstream packets to prevent eavesdropping. The upstream sharing is performed by time synchronising the ONTs and performing Time Division Multiple Access to share the upstream capacity. The NBN in Australia appears to be using Alcatel Lucent's GPON passive optical network, which Alcatel says has a 2.4Gbps downstream and a 1.2Gbps upstream capacity. The NBN documents appear to indicate that each GPON splitter serves some 32 ONTs. The implication is that if all the ONTs use the service then the 2.4Gbps distribution service is shared across 32 services, providing an average of 75Mbps per service. So while the carrier signal in the fibre runs at gigabit speed, if the service were to be fully subscribed and used at capacity then the equally-shared capacity of the network would provide a sustainable service of 75Mbps per end user. And given the typical performance of shared common bus systems it is reasonable to expect that individual subscribers could expect to be connected to a service that operated at level that was equivalent to a 100Mbps dedicated point-to-point service.

So what's the better approach? FttN or FttH?

The FttN would conventionally require some form of active node unit with a VLAN switch, connecting each service via a L2 VLAN to the exchange. There are two critical parts of the FttN architecture: the capacity of the fibre trunk circuit from the exchange to the node, and the capacity of the copper tail loops. The DSLAM function now sits in the node point, and it would be probable that all services provisioned from that node would operate the same DSL protocol. Given the rather severe limitations of VDSL, and the vagaries of vectored DSL, it seems that this model would not go much further than ADSL2+ in the coming years. What does this imply for a network that already supports ADSL2+ tail loops in its copper infrastructure? As pointed out above, what it could possibly do is increase the likelihood of being able to operate the copper pair at full ADSL2+ speed. Whether that's worth the \$30B price tag to do this for everyone in the country is something the Australian voters will need to ponder. The incremental upgrade path of this model would appear to point toward a AON FttH model, where individual customers could replace their copper loops with fibre, and use an active switching element in the node to direct and receive their traffic over a fibre tail rather than over a copper pair DSL service.

The GPON architecture of the FttH service is not without its own problems. The issue here is that all the tails receive all of the aggregate downstream traffic, so that if you are the only user on your GPON splitter then your service will be good, but as your neighbours take up the service everyone takes an incremental hit. Upgrading the base speed of a GPON system requires concerted activity. Not only do you need to up the clock speed at the exchange, but you need to up the clock speed of each and every ONT. The other upgrade option is keep the data speed constant, but divide the GPON splitter into two splitters, and thereby increase the effective capacity per subscriber by halving the number of subscribers per splitter.

In thinking about this, neither of these Ftt approaches really impresses me. The FttN approach is likely to get stuck on ADSL2+, and the incremental upgrade path thereafter involves some rather convoluted node surgery to provision AON services into the nodes. It's likely that a small number of services would be upgraded, and this would require the curb cabinet to be equipped with an active optical unit as well as a DSLAM. The approach was also drawn some criticism in terms of reliability, based on the BT experience with a FttN deployment. [<http://delimiter.com.au/2012/04/30/fttn-a-huge-mistake-says-ex-bt-cto/>] However, what bothers me more is the observation that the last time we tried FttN with the RIM architecture we ended up not upgrading the deployed base at all, and left the population who lived behind these RIMs effectively stranded in digital hell. The most likely fate of this massive expenditure is that it takes a copper network that operates at speeds of at most 24Mbps and offers the consumer a copper network that operates at speeds of at most... 24Mbps. Yes, there is the prospect of VDSL, but the experience with VDSL has been that the adjustment from a 2.2MHz power spectrum bandwidth to a 30Mhz bandwidth creates a huge crosstalk problem, and in the low density suburban environment of Australian cities the node-to-premises copper lengths tend to be longer and not shorter. And yes, I'm somewhat cynical about the benefits of the Vectored DSL band aid. And all to lift the achievable speed to 100Mbps. Where's the next increment of bandwidth going to come from? How could we double the capacity? At some stage the physics of using copper pairs in a harsh external environment simply become too expensive to overcome and we need to head back to fibre optic cables to find a true broadband solution. So in looking at a second round of a FttN approach, I'd rather that we did not head back down the path of making short term cost savings at the expense of future capability.

But I also have concerns with a FttH architecture that relies on a passive splitter approach. Its capability to deliver high speed to individual service points is dependant on the extent to which others who are also residing on the same splitter are also driving their service at speed. While the network was largely TCP rate-controlled Internet this perhaps was not so much of a concern. The services rate adapted to the capacity that was available, and, as in many other forms of sharing behaviour, it was often more effective for 10 users to share a common pool of 100 units of common capacity than it was for the same 10 users to each have exclusive use of just 10 units of capacity. But in an Internet that is

dense with non-adaptive UDP video flows then the benefit of this form of capacity pooling tend to dwindle.

If there was an option for an Active Optical Network in a FttH framework then I think I'd prefer to head in that direction. That's in spite of the considerations of the reliability issues associated with the deployment of active electronics in the node. This approach offers a direct path to increase the capacity of the trunk fibre from the exchange to the node, and a means of increasing the individual capacity from the node to each ONT on a service-by-service basis if need be. In the trunk networking environment in Australia we have already seen the long haul fibre network that was constructed in the mid nineties be upgraded from the original 500Mbps capacity to multi-gigabit capacity through the retro-fitting of DWDM optics, using the original glass. While the edge cables in a FttH environment might not be the subject of such intense levels of capital investment, the reassuring thought is that the megabit speeds achieved through the FttH network are an artefact of the electronics of the system rather than a physical limitation of the cable plant, and there is an progressive upgrade path that does not involve a complete replacement of the edge cable system.

However, an AON FttH network is not one of the options for Australia's NBN right now. But even with what's on offer, making a choice between a GPON FttH and a ADSL2+ FttN network does not seem to be all that difficult. I'd rather do the field work now to replace the copper tail loops in one pass and come out of this exercise with an all-optical reticulation network now. The electronics at either end of the fibre can be replaced at any time in the future without the attendant cost of the field work to replace all the physical copper cable infrastructure. I would much rather see an network infrastructure in place that has the potential to address future needs in communications than spend all this money to provide a national communications network infrastructure that is only capable of meeting today's needs, without any realistic attention being paid to where all this massive investment in silicon is heading. It's like looking at the national transport infrastructure at the start of the twentieth century and deciding to improve the lock system on the canals without paying any attention to the emerging needs of those annoying horseless carriages!

Irrespective of my own personal preferences, I must admit that I find it rather surprising to see some of the subtleties of the differences between FttH and FttN being considered in the mainstream Australian media. The result of a couple of days of media attention on this topic is that we've all had a rapid education in the salient features of advanced high speed network architectures and their relative merits and relative costs. It prompts me to wonder if somehow we could make the support of a massive exercise of IPv6 deployment across the entire national network into an equally prominent election issue. Now there's a thought!

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

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