

Technology Considerations for the Australian National Broadband Network (NBN)

A submission to the Senate Select Committee on the
National Broadband Network

31st January 2014

Dr Craig Watkins
Ing. Kelvin Lillingstone-Hall

Reviewers: **Peter Hitchiner** and **Dermot Cox**

As professional engineers, we have addressed the task of providing input on the NBN through taking a holistic engineering perspective to an important national business problem with a variety of potential technical solutions. We have attempted to distil the essence of technical requirements from the broad spectrum of expectations of the NBN, and apply engineering rationality to identify the most viable and prudent technology choices.

A National Broadband Network initiative is both ambitious and vital for Australia. Experience has proven that a successful nation-building endeavour can only be achieved via careful analysis and input at many levels of the decision-making pyramid. This is especially crucial for the high-level technology selection considerations related to the provision of the premises access network connection.

Preamble

Australia has been conscientiously engaged in the process of delivering a National Broadband Network for at least the last 4 years (and arguably significantly longer). NBN Co has accomplished a great deal since its corporate birth on the 9th of August 2009. The task is viewed by many as the largest infrastructure project ever undertaken in Australia. While there is a very small percentage of the population that question the need for any such undertaking, the overwhelming consensus is that a major upgrade in communications infrastructure is perhaps long overdue and most certainly welcome. We exist in a period of rapid global change, with modern flexible and functional communications capability at the core of our economic and social life. The NBN vision is to supply the network our nation needs today and well into the decades ahead.

A tremendous amount has been achieved by NBN Co and cooperating commercial entities to date. However, the provision of communications infrastructure of the nature and scale of the NBN is an extremely complex undertaking. An opportunity for strategic review has arisen following the change in federal government in September 2013. The first significant public release of information from the NBN Co review occurred on the 12th of December 2013 (with substantial redacted components)^[1]. This document joins other major documents such as the 2010 NBN Implementation Study Report^[2] and an Analysys Mason report from 2012^[3] in attempting to uncover the core requirements the NBN build must meet, and in attempting to detail how such projected capability demand might best be met (by analysis of various high-level options).

The complexity of the strategic review task is clear to all with more than a cursory interest in the NBN. The report released on the 12th of December^[1] reinforces that there is no “magic bullet” and the question of how best to provide the nation's future communications capability is far from settled. However, the last few years have provided clarity on many issues, both from reflection on progress and lessons learned from the NBN Co roll-out to date, and from observation of the broader global industry.

The present contribution aims to add to the national debate by picking up on publicly available information and contributing technical insight in areas that appear to be complementary to analyses already performed either internally or externally to NBN Co. The NBN continues to attract large volumes of comment in a multitude of public fora. Unfortunately there is so much ill-informed comment in public spheres that it is a very difficult process to uncover thoughtful and meaningful

input. Vested interests of many individual contributors is also a readily observable issue. **As professional engineers, our contributions find strong guidance in the Code of Ethics that all members of the Engineers Australia organisation are bound to adhere to. Vested interests have no part to play.** Perhaps more crucially, the strength of comments, recommendations, and any conclusions are carefully noted as appropriate.

Review of the NBN is clearly a complex undertaking, and a truly independent analysis has high value. The limitation to publicly available information may at times be significant, but conclusions and recommendations within this report are made based on broad input, and where not clear from context, the level of firmness of such conclusions is noted. At various points more speculation is called for. The professional engineering tendency is to avoid speculation and focus on hard realities, but thoughtful speculation has a clear role to play when dealing with technologies beyond a 2-3 year horizon.

Our approach is to apply rational engineering analysis to the NBN issue, attempting to identify network requirements and implementation options. Through detailed engineering consideration on matters such as the all-important topic of network contention, we are able to draw a number of solid conclusions and recommendations. Ultimately we believe that this work should inspire further in-depth investigation, possibly within NBN Co, revealing an optimal NBN outcome for the nation.

Many individuals have participated in discussions that have informed this contribution. A number of dedicated individuals have helped prepare, edit, and review, the document. All share the passion and vision for a National Broadband Network. The risks to the nation of taking a wrong step appear very real. Detailed and thoughtful analysis is required to minimise such risks. We intend this contribution to help foster targeted national debate to ensure that as a nation we create an optimised National Broadband Network. **We can not afford to take a wrong step at this crucial point in the nation's history.**

A National Broadband Network is not a panacea. The NBN alone can not ensure that the nation's economy evolves to meet global competitive threats. It can not establish and expand a robust culture of innovation across all sectors of endeavour and it will certainly not magically solve global warming! However, the NBN facilitates such crucial objectives to a greater or lesser extent in each particular case. Within this contribution there is the germ of the prospect of finding a middle-ground solution that satisfies what has been a polarised debate on technology ideology, and in political terms, over recent years. If our nation can work together to maximise potential outcomes on the NBN, then the possibility exists that we can tackle other issues of significance with an equal chance of success.

A first-class NBN outcome that we can be proud of on the global stage is not produced by over spending. It is also not produced by picking what we believe to be a cheaper short-term solution without careful consideration of real costs and real benefits matched with realistic demand projections. It can be provided through intelligent consideration, perhaps some innovative thinking, and a shared vision and commitment. Working together we can ensure our nation creates maximum value by establishing network infrastructure that will serve the country for decades.

A glossary of abbreviated terms is provided at the end of the document.

Table of Contents

Preamble.....	2
Core Recommendations.....	5
1. Executive Summary.....	6
1.1 Key Ideas and Considerations.....	7
1.2 Expansion of Core Considerations.....	9
2. Background Context of the NBN.....	13
3. The Bandwidth Equation.....	16
3.1 Historical Bandwidth Trends.....	17
3.2 Future Application Demand.....	20
3.2.1 The 8K Prospect?.....	23
3.2.2 3D Television?.....	24
3.3 The Evolution of Video Conferencing Technology.....	25
3.4 Video Upload Requirements.....	26
3.5 Small Business Upload Requirements.....	26
4. The Contentious Issue of Contention.....	27
4.1 NBN Co Satellites and Contention.....	27
4.2 NBN Co Fixed Wireless and Contention.....	28
4.3 QoS (Quality of Service).....	31
5. NBN HFC Incorporation.....	34
5.1 The Immediate Future and Beyond.....	36
5.2 Some of the Many Questions for HFC.....	38
5.3 HFC Black-spot Infill.....	38
6. Preliminary FTTdp Discussion.....	40
6.1 The G.fast Promise.....	40
6.2 The Cost of FTTdp Deployment.....	42
6.3 Additional Cost-Benefit Considerations.....	46
7. Large and Small MDUs.....	48
7.1 Small MDUs.....	49
7.2 Mobilisation of a Contractor Ecosystem for MDUs.....	50
7.3 MDUs in the HFC Envelope.....	51
8. Mobile Connectivity.....	52
8.1 Exploitation of NBN Assets for Mobile Augmentation.....	52
8.2 Caravans and Motor-homes.....	53
9. The NBN Deployment Prioritisation Issue.....	54
10. Conclusion.....	55
Appendix A: Contention Basics.....	57
A.1 Shared Links.....	58
A.2 Statistical Usage Patterns.....	58
A.3 The Statistics of 1 Gbps FTTP.....	59
A.4 Contention Issues with a FTTdp Network Solution.....	60
A.5 FTTN Contention Considerations.....	60
References.....	62
Glossary.....	64
Report Authors.....	67
Reviewers.....	68

Core Recommendations

1. An honest appraisal of bandwidth demand trends is required. Objective analysis appears to suggest the risk is very real that demands may exceed the capabilities of a core FTTN (Fibre-to-the-Node) approach within a short time frame.
2. FTTdp (Fibre-to-the-Distribution-Point) must be considered in detail as a mainstream NBN deployment approach. It provides vastly superior network capability outcomes to FTTN, including significant flexibility to meet demand variations, and is likely to be cost competitive with FTTN.
3. Full exploitation of HFC (Hybrid Fibre Coax) capability is recommended, but this requires careful and very detailed engineering planning. Creation of a fully developed long-term technology roadmap for HFC areas must be commenced immediately.
4. MDUs (Multiple Dwelling Units) should be provided with autonomy, stimulating an ecosystem of competitive independent contractors.
5. NBN certified mini/micro/basement DSLAM equipment, together with lead-in and building fibre cabling, could be installed by NBN certified independent contractors and technicians sponsored by RSPs, stimulating a rich ecosystem of installers.
6. A multi-interest coordination committee should be established to ensure NBN build priorities are optimal to meet customer needs and are not subject to wild fluctuation. Full NBN deployment understanding by all parties will enable minor network investment work by current broadband service providers, to service customers where the NBN is some time away from delivery.
7. Fixed wireless should only be deployed in areas where VDSL implementations are impractical, due to the cost/benefit advantage of VDSL.

1. Executive Summary

The nation lacks a flawless crystal ball to project with certainty the bandwidth demands to be met by the NBN over the next several years. We have witnessed a prolonged period of polarised debate with one extreme suggesting 1 Gbps connectivity for the masses is a requirement as soon as 2021, and the other extreme stating first 12 Mbps is enough, then 25 Mbps, and now perhaps 50 Mbps. Amid such divergent viewpoints it is important to agree that flexibility to meet uncertain future demand must be a key factor rolled into any NBN-related technology decision. It is crucial to realise that bandwidth demands from user to user are far from homogeneous. An understanding of average demand is vital for resource provisioning in the backbone network, while appreciation of demand profiles across users is important on the access network side. In both cases, demand evolves in a complex way responding to the entire global ecosystem of technology progress.

- **FTTP** (Fibre-to-the-Premises) clearly provides the flexibility to meet any future demand, but if this is too costly in time and budgetary terms then other technology options must be considered.^A
- **FTTN** (Fibre-to-the-Node) has received a significant amount of renewed interest recently, and some issues related to FTTN deployment will be discussed within this document.
- **FTTdp** (Fibre-to-the-Distribution-Point) is the other natural candidate that sits somewhere between FTTN and FTTP. FTTdp, with average copper lead-in length of around 30m, will be discussed extensively in this report.

A primary conclusion is that FTTdp (in the form of Fibre-to-the-lead-in-pit at the front fence) must be analysed in detail as a potential technology for mainstream NBN deployment.

The recently released NBN Co Strategic Review report^[1] points to average overseas FTTP deployment costs that would be very favourable to full FTTP deployment if such costs were proven to translate to the Australian industry setting. Given indications that NBN Co FTTP costs are substantially above such baselines, it is clear that other technology options must be considered. Globally, FTTN is typically deployed by incumbent operators seeking to maximise the value of their copper assets. The fact that the NBN Co context differs slightly from this norm should not unduly concern us if the cost-benefit relationship is strong enough. Through a smattering of field trials we can extract a reliable indication of the capability provided by a FTTN Vectored VDSL2 deployment in the context of the Australian copper network. The missing link is how this capability matches with the unknown network demand that will emerge over the coming years. The FTTdp alternative provides a far stronger match between capability and the unknown demand and, perhaps surprisingly, appears to do so on a cost-competitive basis to FTTN.

A The NBN Co Strategic Review document indicates that FTTP may be too costly, but there is reason to suggest further attention to business practices, implementation matters, and high-level network architecture, can still hope to achieve significant cost reduction. Additionally, whether an option is too costly or not requires a full assessment of both costs and benefits, perhaps coupled with an analysis of the nation's ability to foot the bill.

1.1 Key Ideas and Considerations

The remainder of the executive summary section highlights key dot points and then expands on more major topics below.

- **FTTN/FTTB appears logical for both small and large MDUs**
 - Small MDUs (Multiple Dwelling Units) may benefit from micro-node equipment
- **Exchange-based VDSL2 deployment in many smaller regional towns offers significantly greater capability than fixed wireless**
 - Little expectation of improved outcomes can be provided for the more dispersed percentages of the population, but the contention aspect of fixed wireless makes it a largely retrograde step for smaller communities densely clustered around existing exchange facilities
- **FTTN presents a major risk for mainstream deployment in terms of the probability of rapid bandwidth demand growth**
 - Consideration of bandwidth demand for Ultra HD 4K video provides a worthwhile proxy for anticipation of broader demand over the next few years
 - Flexible provision of variable capability from user to user is comparatively limited in FTTN compared to FTTP or FTTdp
 - VDSL profile selection locks the global network configuration with FTTN
 - Migration to FTTP or FTTdp may need to occur on such a short time scale that initial FTTN deployment becomes a clear misstep in hindsight
- **FTTdp must be thoroughly considered as eliminating some of the major risk factors of FTTN and providing improved deployment prospects compared to FTTP**
 - Relatively short copper length of FTTdp and the likely availability of bonding and phantom mode for drop pairs ensure high VDSL2 bandwidth (up to 250 Mbps) and flexible upgrade to G.fast and FTTP on an as-needed basis
 - FTTdp shares with FTTP the prospect of a single network build process
 - FTTdp eliminates large street cabinets of FTTN and has the potential for significantly smaller street impact than FTTP
 - The micro-node implementation approach of FTTdp provides some flexibility in VDSL profile selection in addition to reducing the impact of the profile selection trade-off due to higher overall data bandwidths
 - Careful analysis may reveal significant deployment cost savings on the Local Fibre Network (LFN) side of the LNDN (Local Network Distribution Network) compared to FTTP due to flexibilities provided by lower FTTdp LFN fibre counts

- **HFC incorporation makes sense, but upgrade paths must be carefully analysed**
 - HFC (Hybrid Fibre Coax) is originally a broadcast mode network for supply of a few video channels while the major user trends are toward selection of many individual video and data channels from millions of worldwide sources, as well as greater use of two-way peer-peer capacities. Technologies such as GPON (Gigabit Passive Optical Network), FTTdp, and FTTB (Fibre-to-the-Building/Basement) are well-suited to this usage mode. DOCSIS (Data Over Cable Service Interface Specification) 3.0/3.1 are promising, but is it enough to overcome the basic broadcast network structure issues?
 - Careful consideration must be given to the 'end goal' in HFC areas. Is this RFoG EPON (Radio Frequency over Glass Ethernet Passive Optical Network), ultimately being a form of FTTP, or does node splitting continue to occur on an ad-hoc basis as bandwidth demand rises?
 - Does a FTTLA (Fibre-to-the-Last-Active/Amplifier) upgrade process suffice to ensure that fibre is gradually pushed deeper into the network, providing support for future demand?
 - HFC is noted to be relatively power hungry. How can this be reduced over time?
- **A multi-interest coordination committee to oversee NBN deployments must be considered**
 - A committee is far from the only model that might achieve the desired governance outcomes, but considering this option is likely to clarify transparency, stability and performance objectives of the deployment prioritisation issue. The intent is to make overall broadband services provided by all current broadband providers more effective to meet customer needs during the lengthy NBN build stage as well as assist NBN Co to stage the NBN roll out to meet the most urgent needs first (after first determining how urgency should be defined and assessed in this context).
- **A healthy ecosystem of independent contractors for MDUs should be encouraged**
 - NBN Co certified mini/minor/basement DSLAM (Digital Subscriber Line Access Multiplexer) equipment, lead-in fibre cabling and building fibre cabling could be employed at the customer end of the NBN and these could be installed by NBN Co certified independent contractors/technicians sponsored by RSPs (Retail Service Providers), etc.

1.2 Expansion of Core Considerations

While substantial uncertainty exists in the precise capability demanded of the NBN access network in future years, it is clear that bandwidth demand growth shows no signs of slowing. Exchange hosted ADSL2+ service capability broadly lags demand by a significant margin at present. HFC networks fare better currently, but the need for dramatic capability improvement to meet future demands is clear.

The vision of a National Broadband Network serving the needs of the nation for many decades is one that holds strong consensus support.

Bandwidth demand from user to user has historically shown tremendous variability. The popularity of Facebook and the trend toward ever richer web content ensures that 'light' users will experience massive continued growth in demand. The emergence of ever more high bandwidth consuming applications translates to a growing gap between average user demand and the demand of 'heavy' users. The NBN must meet the needs of all users as effectively as possible.

Flexibility is required to meet the varying demands of users. FTTP is able to meet all demands with the possible criticism of provision of more capacity than required by most users for the immediate future. Other options such as FTTN and FTTdp must be assessed with flexibility of user demand in mind.

FTTN makes plenty of sense in the MDU (FTTB) setting where copper lengths are short and the practical difficulties in fibre cabling for FTTP are substantial. However, there are very real limitations of the FTTB Vectored VDSL2 approach which should not be overlooked. Short copper lengths may provide great network bandwidth for MDUs, but the need for agreement on a single VDSL profile limits utility (to an unclear extent). The prospect for upgrade to G.fast speeds in the large MDU environment with crosstalk and potential susceptibility to radio interference also needs further investigation. The higher transmission frequencies employed by G.fast make it especially sensitive to crosstalk and interference from broadcast radio transmitters. A single crosstalk-inducing active pair can reduce G.fast bandwidths from 500 Mbps over 100m to 60 Mbps^[6]. Advanced vectoring techniques have been observed to restore the bulk of the single-line throughput, but this is recent research output, is extremely CPU and power intensive, and is not a tried and tested commercial product offering at this stage. The restriction to a single crosstalk-inducing pair is also not in any way representative of the projected future usage of G.fast in the MDU environment.

FTTN using small micro-node devices should also be considered in the small MDU setting where the alternative is time-consuming fibre cabling within the complex. Should a FTTN deployment occur in the area of such a small MDU, then careful consideration will be required for where such MDUs are served from the main node or from a smaller node within or adjacent to the MDU (most

obviously in the MDU's existing Telstra MDF). The prevalence of shared copper lead-in introducing significant crosstalk must be assessed carefully. Admittedly the possibility of significant crosstalk is also a consideration for many SDUs in potential FTTN deployment scenarios.

Recognising the fact that there is no "one size fits all" solution, and the uncertainty with regard to future demand projections, it appears sensible to ensure that MDUs are given the opportunity to provide their own fibre cabling and opt for FTTP at the time of NBN Co connection (if at all feasible). The cost to the body corporate of such an approach will certainly dissuade many, but the cost to NBN Co would perhaps be expected to be on par with the FTTB option. Operating overlapping FTTP and FTTN networks in the same geographical area is a complication that is likely to start appearing almost immediately in an FTTN deployment scenario due to the "fibre on demand" model necessarily introduced to provide capacity to higher demand users such as small businesses in largely residential (and hence FTTN) areas.

Due to short copper lengths, MDUs of all sizes make sense for FTTB deployment based on vectored VDSL2. Flexibility dictates user-pays "fibre on demand" should be provided where possible. This implies an initial FTTP build option for those willing to fully fund the fibre on the downstream side of the splitter.

Should VDSL be deployed in any NBN context, it appears crucial to ensure that the highest level of hardware capability is present, especially given the high level of uncertainty with regard to demand projections over as little as three to five years. Vectored VDSL2 with capability for bonding and phantom mode must be considered from the outset. The reality of the FTTN approach suggests major limitation in the prospects for bonding and phantom mode due to pair shortages in the copper bundles exiting the node. This is a very serious limit for the longevity of a FTTN 'solution'. We note that FTTdp on the other hand, while exhibiting far less gain from vectoring due to the relative absence of shared lead-in cabling, is likely to benefit significantly from bonding and phantom mode. In fact VDSL2 without vectoring would generally work well in the FTTdp situation, and bonding and phantom mode produce substantial VDSL2 capability over short copper lengths. 250 Mbps is widely reported in such a setting. Importantly this is also the 'ideal' deployment scenario for G.fast, should such a future upgrade become necessary.

Exchange-based VDSL2 looks to provide vastly more capability for many smaller regional towns (with pronounced clustering of premises around centralised exchange locations) compared to fixed wireless. As user demand behaviour patterns continue to evolve, RF contention in fixed wireless areas is likely to become a major issue. An exchange-based VDSL2 implementation provides significant flexibility to many regional towns, with greater depth of fibre deployment possible over time on an as-demanded basis. We must acknowledge the fact that network expectations being set by early IP TV services delivered over ADSL have severe consequences for fixed wireless. Even without this threat the general increase in prevalence of streaming content is cause for major alarm. Over time the video quality expectation is sure to rise, pushing up streaming data rates. We can also expect IP TV service market penetration rates to increase. Constant data streaming at high bit rates over fixed wireless from multiple users in the one sector is one situation where fixed wireless capacity limitations will be obvious to all users.

Many areas where an exchange-based VDSL implementation is likely to be beneficial are already served by ADSL. These areas can expect very poor subscription to fixed wireless NBN Co offerings, in spite of any ignorance of the RF contention issue and the relative attractiveness of 25/5 Mbps headline rate service offerings compared to ADSL. Early fixed wireless users are noting a huge improvement in upload speed compared to ADSL services, and there is presently good early customer feedback helping to encourage NBN Co service uptake. In spite of this, the prognosis for service uptake in areas well served by ADSL is not strong.

We often hear the 7% slated for satellite and fixed wireless associated with the word 'remote' and the dispersed nature of the population across wide areas of the nation. A proportion of this 7% lives in regional areas but in densities similar to suburban settings. Many such population clusters are very close to larger regional centres.

A significant proportion of the 7% of the population currently slated for fixed wireless and satellite connection are unlikely to have any realistic alternative. Where opportunities exist to provide vastly improved network capability for minimal cost, they must be pursued. Reducing to even 6% the number of people on the wrong side of the NBN 'digital divide' is a major win for the nation.

Predicting the level of general broadband demand in the future is clearly problematic, but very few voices would today claim that 25 Mbps is enough in the way some may have considered it to be only a few short years ago. The spectre of Ultra HD video is now upon us, and while video is not the only application of interest to the NBN, it does provide a convenient analysis point. Netflix provides a baseline reference point for 4K streaming at 15.6 Mbps, but we must conservatively assume that a single 4K stream requires on the order of 20 Mbps to 25 Mbps. Multi-person households also face a reality of multiple simultaneous users of the access network in a way that was simply speculation a few years previously. While it is still difficult to forecast how an average user might employ a 1 Gbps connection, it takes very little imagination to see how 100 Mbps might be used by an 'average' user in a very small number of years. The conclusion is that our immediate target bandwidth provision needs to be significantly higher than this 100 Mbps figure, or at the very least the network implementation needs to have ready flexibility for simple on-demand upgrade.

Widespread deployment of FTTN runs a very real risk of supplying inadequate network capability. FTTP on the other hand appears to have failed to meet most promises of deployment cost and timetable. **FTTd can be viewed as a middle-ground solution for brownfield NBN deployment.** (For greenfields there is strong reason to suggest that FTTP is the only sensible choice.) The very short copper lengths of FTTd imply that VDSL rates will be high enough to meet the needs of all but the most demanding users for the immediate future. G.fast also promises a convenient upgrade path. An on-demand, user-pays fibre drop model completes the equation in terms of flexibility and network longevity. The small node devices are powered from the premises over the copper, and there is no need for large (unpopular) street cabinets. The likely availability of multiple lead-in pairs for many premises allows pair bonding and phantom mode for maximal VDSL2 (or G.fast) performance. Splitter fan-out is significantly reduced compared to FTTP, allowing the possibility of street furniture greatly reduced in scale, including solutions where the

splitter frame is retracted into a pit. Greatly reduced LFN fibre counts compared to the default FTTP model used by NBN Co to date introduces the prospect of substantial network deployment cost savings on the LFN side in addition to major savings on the customer connection.

Mainstream FTTdp NBN deployment must be thoroughly analysed. In contrast to FTTN it promises a single network build, long-term solution. The fibre drop portion of FTTP is eliminated, translating to a substantial deployment cost saving.

2. Background Context of the NBN

It may be observed that FTTP (or a broad equivalent) has been promised for much of the last 25 years. Under the direction of the previous (2007–2013) Australian government, NBN Co was charged with the challenge of attempting to deliver FTTP to 93% of the nation's premises. Such a task has been widely claimed as Australia's largest ever infrastructure project, and recognised as ambitious compared to many other overseas deployments.

Nonetheless, a panel of experts concluded that the nation needed a major boost in national communications infrastructure to facilitate emerging economic activity. Communications technology is at the very heart of our modern society and grows steadily in importance as the years progress. We rely on ubiquitous network access like never before. A healthy telecommunications ecosystem enables leading-edge service improvements on a consistent basis. There are perhaps a multitude of reasons for the historians to argue as to why this appears to have not proven to be the case in Australia. However, the tyranny of distance clearly has a significant role to play domestically.

The conclusion that Australia's communications capability requires a major boost is as strong or stronger today than when NBN Co was formed.

Demand growth for broadband communications has been promising to exceed the capabilities of copper pairs for decades. ADSL was the first technology to provide a significant increase in copper pair utility in a commercially viable way for the masses. The later emergence of VDSL did not have a broad impact initially due to a higher implementation complexity (cost), and the marginal value over ADSL2+ being small for most users. Indeed, for larger copper loop lengths (those premises more distant from exchange sites), exchange-hosted VDSL offers no improvement over ADSL.

VDSL does offer significant capability advantages today for more moderate copper loop lengths (up to 500m or even 800m range, although optimal performance comes at far shorter lengths). The implementation complexity is no longer a significant barrier to deployment. Partially this is due to steady improvements in processing capability and the emergence of tailored silicon technology implementations (Application Specific Integrated Circuits – ASICs), in turn facilitated by economies of deployment scale.

While ADSL and the capabilities of DOCSIS (Data Over Cable Service Interface Specification – the standard for two-way data communication in the coaxial cable network world) have delayed the need for taking the step improvement in network capability that only fibre can provide, rapid and prolonged increases in network demands make this step seemingly inevitable. It appears that the conclusion of inevitability was precisely that drawn by the panel of experts leading to the formal creation of NBN Co in August 2009.

We must ask whether subsequently anything has changed. Is FTTP still viewed as inevitable, and if so, has our assessment of the timeline changed in any significant way? The answer is that DSL technology has continued to push closer to the limits of what capability can be extracted from copper pairs. (DOCSIS has made parallel improvements in data capacity provision over coaxial cable networks.) VDSL technology is now mature, and with bonding, vectoring, and phantom mode, is reported to achieve up to 250 Mbps for short copper lengths. G.fast utilises additional transmission bandwidth and some more novel sophisticated processing to push the capability of

copper pairs to 500 Mbps and potentially significantly further. G.fast over twisted pairs promises 1 Gbps and more with future iterations of the standard. For the cable world, Huawei has recently announced G.fast technology that provides a similar data capability without impact on legacy cable services^[8].

The conclusion that fibre is the end-game is no longer as clear-cut as it may have once seemed.

FTTP makes undeniable sense for new development areas (greenfields). It is also the only sensible choice where it can be deployed economically, even if this is at a marginal additional expense to other options. The flexibility of capability afforded by a full fibre solution is robust. There is no serious expectation that fibre's capability will be bypassed by other technologies such as wireless. Wireless capability receives a significant amount of media coverage, but largely the fundamentals are poorly understood by those suggesting it is a real alternative for the fixed-line network. However, DSL technology has proven capable of providing all but the more extreme demands for the foreseeable future on very short copper loops.

FTTdp with very short copper lengths (perhaps 30m average) offers a level of capability that removes the need for widespread premises fibre connections for conceivably longer than a decade or two. Given the substantial implementation cost (and delivery time) consequences of fibre drop deployment at scale, the option of taking fibre only as far as the Distribution Point must be properly assessed. Construction of a FTTdp network is still a “do it once, do it right” option, as upgrade to FTTP can easily be accommodated on an individual demand basis as required. There is little to suggest that bandwidth demand of an average user will exceed that provided by short-loop FTTdp employing VDSL2 or G.fast over the next 10 to 15 years. On the same timetable we can safely suggest that a small but significant percentage of premises will see value in funding a fibre drop. The short distances (largely on the premises site) makes a user-pays on-demand approach to FTTP upgrade viable.

If full FTTP could be provided economically the chances are that the present NBN Co Strategic Review process would never have started. FTTN has been proposed as an alternative for some time, and this option has already been raised to prominence in the initial Strategic Review report. HFC is also back on the NBN Co radar. There is now a compelling case for proper detailed consideration of FTTdp to occur.

Ignoring the issue of excess remediation and maintenance of longer copper loops required for FTTN, while 50 Mbps minimum capability offered is better than that available to many today, it does not provide the long-term solution that FTTdp offers. At the same time, there is a very real chance that the short-term economics of FTTdp will be better than FTTN. We are far from being able to conclude one way or the other on this last point, but initial back-of-the-envelope calculations are very favourable (such calculations are outlined in the relevant FTTdp report section below).

FTTdp (in the form of Fibre to the lead-in pit at the front fence) involves a LNDN (Local Network Distribution Network) build that is very similar to that of an FTTP implementation. The LFN (Local Fibre Network) part of the FTTdp deployment appears to suggest a drastically reduced fibre count compared to the existing NBN Co FTTP model, but is aligned with FTTP deployment options with

dual-stage splitter implementations. There is hence very little contained in a FTTdp deployment that would not be directly applicable to migration to FTTP should this become necessary (under scenarios of continued rapid bandwidth demand increases). The most likely outlook is for a small proportion of FTTP to be deployed over time on an as-demanded basis. FTTdp is flexibly positioned to meet such demand.

FTTN cabinets can supply space for splitter frames to enable on-demand FTTP upgrades. Given the reduced capabilities of FTTN compared to FTTdp, there is likely to be a considerably higher demand for such FTTP upgrade. However, we can expect the cost of a long fibre connection to suppress the latent demand for individual FTTP connections in a major way. Ultimately dissatisfaction with a FTTN solution is likely to grow. Under even moderate assumptions for bandwidth demand growth the outlook for FTTN prior to the need for a subsequent major network build phase is a handful of years at best. The number of years predicted varies depending on the flavour of crystal ball employed, but the risk of an extremely short tenure of FTTN is very real. **We need to consider little more than the reality of multiple device content streaming within premises today to appreciate the high risks involved with mainstream FTTN deployment.** (We elaborate on this point in the following section, but note that to obtain a proper evaluation of this risk requires detailed analysis.) Given that even the FTTN build is projected to take many years, the possibility is strong that the FTTN methodology will be abandoned long before the initial build is complete. This would potentially leave NBN Co struggling with the disaster scenario of refocussing to a likely FTTdp or FTTP build process at the same time as needing to improve capabilities in the areas where FTTN has already been deployed.

If FTTN was already a reality in mass deployment in Australia today, then we would probably still be discussing the need for commencing FTTdp deployment immediately in order to reliably meet network demands in 10 years time.

While FTTN would provide a significant capability increase compared to what we actually possess today, this still falls well short of all but the most incredibly limited projections of future demand on the 5-year-plus horizon. FTTN falls short both with supply of average capability and in flexibly meeting the demands of that small but significant percentage of more demanding users. These more demanding users are likely to correlate strongly with growth areas of the national economy.

FTTN may have a valuable role to play in small regional towns where it supplies far greater capability than fixed wireless. It also has a role to play in the FTTB guise in large MDUs where the cost of fibre cabling is prohibitive. For smaller MDUs a FTTdp model parallels the FTTB approach if fibre is taken to the MDU's copper pair MDF. **It is far less clear that FTTN has a core role to play in the broader NBN deployment once we allow proper consideration of the FTTdp option.**

FTTN is a sensible step to take where we can have strong confidence that it will provide sufficient capability for a minimum of several years, and where the deployment cost and implementation time of other options with improved network capability outcomes are substantially larger. FTTdp promises enormous network utility improvement compared to FTTN, and may provide this at little or no cost increase to a FTTN baseline. Mainstream FTTdp deployment simply can not be overlooked at this crucial stage of the NBN.

3. The Bandwidth Equation

Reliable prediction of future bandwidth demand is far from a trivial task. The 2012 Analysys Mason report^[3] (page 35) suggests that 100 Mbps is required for a single 3D HD TV stream, and 300 Mbps is required for a single 3D Ultra HD TV stream. Such projections are clearly at the extreme high end of bandwidth estimates. A projection at the other extreme end is provided with the December 2013 NBN Co Strategic Review document^[1], starting at page 78. Here it is suggested that 4K HD TV streaming could require as little as 12 Mbps with continued compression improvement to 2023.

The difficulty in determining a precise requirement for video is related to the difficulty in understanding exactly where on the bit-rate versus quality continuum future market demand will lead us. We have a solid understanding of the trade-off for HD video formats (although even here the issue is complex due to differences in video scene activity levels), but the outlook for Ultra HD formats and the promise of 3D television is hazy from the perspective of most observers. Ultra HD is set to become a major force in the market in a very short period of time. We can reliably predict rapid consumer adoption. However the 4K format is not yet familiar enough to understand how much quality we might be prepared to sacrifice to view content at our convenience.

The jury is yet to return a verdict on when/how/if 3D television will become established in a significant way in the market. We would be foolish to discount the prospect of 3D television entirely, but must allow that the future of 3D television could be dramatically different from early commercial 3D TV offerings.

We note that the HEVC (High Efficiency Video Codec) standard^[11] provides profiles to support the high bandwidth claims made in the Analysys Mason report. Indeed, the HEVC standard also supports profiles that exceed the 300 Mbps figure by a very wide margin. However, strong objections can be constructed against the points made and conclusions drawn in both the Analysys Mason report and the NBN Co Strategic Review document. Ultimately the large range of demand uncertainty we are left with ensures that NBN network technology decisions made without further detailed analysis of the bandwidth demand equation would translate to decisions made on the basis of almost complete ignorance. **While we can not have absolute certainty of the future, we must ensure weighty NBN technology decisions are made on the basis of real information and not ignorance.**

We need to look only as far as Netflix to obtain some idea of what a baseline of 4K Ultra HD video implies for data bandwidth. Netflix announced wide availability of headline shows in 4K format at 15.6 Mbps at the recent January 2014 CES trade-show in Las Vegas^[14]. It would however be wrong in a very big way to suggest that 4K resolution is fully served by 15.6 Mbps. The trade-off between bit rate and quality will be discussed below, and this baseline of 15.6 Mbps will be put into proper perspective.

Ballpark figures bracketing future demands are a crucial element of the decision making process with regard to selection of NBN policies and priorities. The “bandwidth equation” must thus be analysed in sufficient detail. We start by considering historical trends in the provision of access technology bandwidth, and examine the issue of video bandwidth in some detail. **The NBN is vital for more than movies and video conferencing, but video bandwidth provides a reliable indicator of mainstream bandwidth demand from an access network perspective.**

3.1 Historical Bandwidth Trends




Figure 1: Image Credit -- NBN Co

Many industry and research observers have replicated historical plots of consumer access technology bandwidths, showing an exponential trend (linear trend on log scale). Projected a number of years into the future these plots appear to show a demand for gigabit access speeds as soon as 2021. While not summarily dismissing historical trends, we must consider the validity of using past data to extrapolate future demand. (Figure 1 presents an image extracted from NBN Co documents. Many similar diagrams can be located on the web.)

Of prime significance is the fact that the “Moore's Law” trends of the past are most obviously a combination of advances in the enabling access technology (in turn linked to advances in cheap, low-power computing technology) and inherent consumer demand. The equation is far from simple, and it is vital to note that an entire ecosystem of high technology products and services underlies what might be considered to be inherent consumer demand.


Quotes such as “14.4 kbps ought to be enough for anyone” abound, but hide the reality that early access network technology (dial-up modems in telephony bandwidth) were invariably a limiting bottleneck. While PCs in the early 1980's achieved the giddy speeds of 1 MHz CPU clock with 1Mbyte RAM and 20 kbyte of hard disk, 300 baud modem speeds were already felt as a significant constraint by users even though the only 'applications' were remote login to 'mainframe' computers or university campus computer networks, and BBS systems (bulletin board systems).

Consumers quickly adopted higher speed 1200 bps, 2.4 kbps, 14.4 kbps, and 19.2 kbps modems as unit prices fell and installed ports increased in the burgeoning dial-up market. By the mid 90's 'blistering' speeds of 56 kbps were available through a large number of modem manufacturers and unit pricing was such that mass market adoption occurred quickly in consort with the blossoming of the World Wide Web.

Research attention had already turned to escaping the bounds of the telephony bandwidth restriction in the early 90's. What became the widespread ADSL technology was put through its paces in a number of labs around the world while standards were debated. The core concept was that the

prevalence of fibre in the PSTN networks implied that limitation of a single connection to 64 kbps in the ISDN/digital telephony network paradigm present up to that point, was unnecessary. Using a higher analogue bandwidth on the copper loop (last mile from the exchange or local office), could supply significantly greater data rates. The bandwidth exploited for data transmission was so much greater than that used for voice telephony that the impost related to separating voice and data services on the same pair through simple filtering devices was minimal. Consequently ADSL technology allowed users to have an “always on” connection without interfering with other voice telephony usage within the premises.

In the present ADSL2+ incantation, speeds of up to 20 Mbps down and 1 Mbps up are achievable for those close to an exchange. For those premises more distant from the exchange the achievable speeds drop quickly. This has been a persistent source of frustration for many consumers for much of the last decade. In more recent years a percentage of consumers have been able to meet their thirst for higher connection speeds through HFC networks in some areas, or via 3G and 4G mobile data networks. Such mobile networks supply data at relatively high cost, with monthly allowances typically in the range of 2 Gbytes, compared to a typical allowance over the 'fixed' network of something like 200 Gbytes.



Source: Finnie (2006)

Figure 2: An alternate projection.

Figure 2 shows an alternate projection to that presented in Figure 1, made some years previously (in 2006). The extrapolation shown appears to predict 250 Mbps today and 1 Gbps by mid 2016. While interpreting the above graph in such a way is likely to be significantly outside the scope of the projection intention, it does highlight an important point. It is easy to choose data points differently and present a steeper trend, seemingly pointing to 1 Gbps demand a lot earlier than 2021.

While a judicious selection of data points can appear to show a broad linear trend on a log-scale plot of access speed versus date, the reality is a complex interplay of restrictions limiting bandwidth supply. Such plots can be accused of showing the supply side of the bandwidth picture and leaving us to guess at implications for demand. A trend of rapid, ostensibly exponential growth, is clear, but direct extrapolation to claims of 1 Gbps demand in 2021 are unjustifiable.

There are many forces in the technology world that could potentially align over the next 7 or 8 years to cause average consumer demand to increase to the 1 Gbps range (although this remains at the extreme high end of expectations). However, the forces behind future bandwidth demand growth are sufficiently different from the historical growth enablers/causes to invalidate a strong conclusion from direct extrapolation of access technology bandwidth historical data.

We must be open to the possibility of continued demand growth that approaches an exponential trend, but also consider that the reality may be a significant slowing of the demand growth curve. There is no reason to suggest a plateau effect, but the possibility exists that a 200 Mbps to 300 Mbps average service level will suffice in 2021. 50 Mbps or 100 Mbps may even suffice, but realistic assessments of demand made today suggest that these rates will quickly be seen as inadequate, well before 2021.

Crucially, bandwidth demand is not uniform from user to user. For backbone network provisioning it is average usage coupled with temporal variations (peak usage considerations) that are important. For the access network component we must account for usage spikes and variation from premises to premises. Ultimately this involves the prospect of technology compromise. FTTP can flexibly cater for usage variations but at criticism of potentially overspending on infrastructure. FTTdp comes close to offering the same level of flexibility afforded by FTTP, while FTTN is more obviously a compromise technology. The question becomes one of whether the service levels offered by FTTN are good enough and whether we evaluate the risk of steady bandwidth demand increases as high or low compared to the capability offered by FTTN. In order to answer this last question we must turn to consideration of future bandwidth demand trends.

Rational selection between various technology options involves analysis of both costs and benefits. A discussion of costs is contained in a later section. The present focus is attempting to determine the network capability benefit provided by relating capability understood to be delivered by particular options, to bandwidth demand considerations.

3.2 Future Application Demand

4K Ultra HD TV penetration rates are likely to be rapid compared to the introduction of HD TV. We must conservatively assume a bandwidth demand of 20 to 25 Mbps for a single 4K data stream, and that a significant amount of content demand will occur over IP networks.

The outlook for 3D television remains unclear. There are reasons to expect 3D television to become of wider market interest with the paradigm of larger 4K screen sizes. The bandwidth implications are unknown but may be significant.

8K Ultra HD has a hazy outlook at present compared to 4K. It provides little additional bandwidth demand consideration to that already implicit with the likely need to stream multiple 4K video feeds into a single (not far from average) premises within a very small number of years.

Access technology (dial-up modems) 20 years ago provided bandwidth where audio streaming was a major challenge, and this became a significant mass-market application as modem data rates increased. Today we face similar challenges with video content. All but the fastest ADSL connections struggle to supply reliable HD video, while Ultra HD video is already clearly on the technology radar. A Sony 140cm (55") 4K TV retails for less than \$4000, and a Hisense 165cm (65") 4K TV can be purchased for less than \$3000. Intense competition in this sector of the technology market will ensure rapid downward price pressure. The level of price reductions seen in recent years for HD TV screens is an indication of what the future may hold for 4K units. Reputable brand 81cm (32") HD TV units can be purchased from less than \$300 today, and 165cm (65") screens today cost a similar amount to the 81cm (32") screens of as little as 5 years ago.

HD television represents a step up in quality from legacy analogue television. Importantly, LCD/LED screen technology has enabled a paradigm shift to much larger screen sizes. This adaption in viewer expectations/demand is steadily occurring and has significant inertia in the complex ecosystem of video content generation (movie and television programming). **The trends are clear, irreversible, and accelerating.**

Analogue television broadcast has been decommissioned. Digital radio TV broadcast is now the reality for those without Cable TV broadcast, and we can expect to see the dominance of full HD digital broadcast increase steadily over time. We must ask whether the radio broadcast paradigm of television is likely to persist with the move to 4K content. The question is not so much whether radio broadcast will continue to exist, but the extent to which the move to IP delivery of video content of all types might be hastened or accelerated by the emergence of 4K. Part of the answer to this question involves an understanding of the transmission bandwidth requirements of Ultra HD 4K content. Improvements in computing resource availability over the last 15 years translates to a significant improvement in achievable video compression performance. Coupled with the inherent flexibility in the Digital Video Broadcast (DVB) standards, this leaves open the possibility of supplying 4K transmission within existing broadcast channel spectrum allocations.

Practical realities of the future of broadcast television are far more complex than this simple equation of compression technology. Expecting rapid adoption of 4K broadcast by the industry is perhaps unrealistic given the multitude of factors at play. The 4K revolution can be expected to be rapid in comparison to the HD revolution, and the effect of alternative delivery mechanisms is likely to be far stronger. The likes of 'new' media entities such as Netflix can not be underestimated^[9]. Such companies have established, and explosively growing, footprints, with an 'immediate' ability to supply new content formats. The ability to monetise the early adopter market has never been greater, and this is sure to drive rapid progress.

We can conclude that 4K Ultra HD adoption is likely to be swift by historical standards, and both enabled and driven through access networks such as the NBN. It is difficult to be definitive in terms of stating a bandwidth requirement, but 25 Mbps is likely to represent a sensible assumption. (4K video content can easily be compressed to 1.5 Mbps, but the quality will be similar to the low quality delivered at 1.5 Mbps today from the likes of SBS On Demand and obviously completely inadequate for display on larger format screens.) The initial Netflix offering is reported to use a video bitrate of 15.6 Mbps, but we must assume the bandwidth-versus-quality equation will migrate upwards slightly over time. Access technology limitations are likely to have contributed in a significant way to the selection of the lower 15.6 Mbps rate.



Figure 3: 4K Ultra HD TV is already here! (Image Credit: Sony)

4K TV screens from major manufacturers include hardware-based implementations of the HEVC (High Efficiency Video Coding) standard. While Netflix makes a strong case that disk-based content distribution methodologies are unlikely to prevail in the 4K paradigm, the reality is that TV screens do not differentiate in terms of what channel/path their video data takes. We can expect healthy global competition to develop in the industry of shifting bits to 4K TV screens. IP network distribution will undoubtedly be a major leading method used to deliver this 4K content. Competition will ensure that higher quality rendered images will predominate over time.

1920x1080 pixel resolution of HD screens are impressive only by historical reference. This resolution is limited by almost all other analyses. An 18cm (7") tablet such as the most recent Google Nexus device, presents crisp and clear video with full HD content. While a jump to 4K

resolution of 3840x2160 pixels on an 18cm (7") screen seems like overkill today, it is most certainly true that full HD resolution even for a 18cm (7") screen is limiting. Text content is most telling for the small font sizes that come from attempting to read sizable documents on a small format device.

The equation is similar from the perspective of something like a 69cm (27") monitor at desktop screen distances. The geometry translates to each pixel taking about four times the visual field size of the pixels in the 18cm (7") device case. We might be happy to read two pages of text on such a monitor, making the pixel to font pitch number roughly similar to the 18cm (7") device example. Ultimately this translates to fuzzy characters at HD resolution, and 4K resolution would have significant advantage.

It is important to note that despite these advantages, there has been no significant industry movement towards higher resolution displays for desktop computers yet. Some Macintosh computer products have led the field slightly, and have had a good following with graphic artists and similar occupations. However the expense in producing both the display itself, and the complexity of the hardware to drive it, coupled with a relative decline in the desktop market with attention to the mobile device sector, has all contributed to this neglect. The larger screens used for TV (165cm or 65" plus) is where the most obvious limitation of HD presents, and this is where we are now seeing the emergence of the 4K paradigm.

The conservative projection is that within 5 to 10 years there is a very real 'risk' that a large number of premises will have peak demand equating to as many 4K streams as there are people in the premises.

We can assume a baseline bandwidth demand of 15.6 Mbps (Netflix initial offering), but must conservatively assume a demand of more like 20 or 25 Mbps for adequate 4K content. Early subjective testing of the HEVC codec for 4K resolution shows only minimal levels of perceived distortion over a collection of 3 five-second test sequences for reducing bit rates^[10]. The test was designed to display the deficiencies of the PSNR objective measurement when comparing the HEVC codec to the older AVC codec at high resolutions. We can only extrapolate the results of such tests with extreme caution to the implication for 4K bandwidth demand over a broad range of content. **The conclusion can only be that 15.6 Mbps must be taken as a baseline and not assumed to be representative of what will be demanded by customers over the next few years.**

Coupled with this fact and the industry inertia resisting standards upgrades once 4K screen units ship in production volumes with hardware HEVC implementations, the suggestion of a 12 Mbps bandwidth requirement as presented in the NBN Co Strategic Review document is out of the range of a reasonable projection. Admittedly, there are still 9 years of technology advances until the 2023 date mentioned. The possibility that sophisticated processing could be implemented in a unit external to legacy screens by that date exists, but must be regarded as an outside chance at best. Even if we are open to such a route to compression technology advances, the reality appears to be the HEVC codec for at least the next decade. **We must conservatively conclude a 4K video data rate requirement of between 20 Mbps and 25 Mbps.** Additionally we must not assume that consumer demand for 4K content will be strictly limited to 25 Mbps. The 25 Mbps figure simply represents a reasonable upper figure on an uncertain projection, and must not be interpreted as a guaranteed hard demand limit. The implication for NBN provisioning is obvious.

It would be remiss not to mention the Google-backed VP9 Codec technology. This has been developed in an open source environment with the intention of providing a royalty-free codec solution that improves on the earlier MPEG video codec offerings. VP9 exploits advances in processor capability to achieve significant gains in output video quality compared to earlier MPEG codecs. At low streaming data rates (suitable for use over mobile cellular networks) the quality improvements compared to the earlier generation of codecs is readily observable. VP9 supports 4K video formats, but it is unclear at this stage how much market penetration will be achieved by VP9 in the 4K space. The backing of Google and the advantage of a royalty-free solution may make it an attractive option for 4K IP network use, including YouTube, despite the significant increase in bit rates required compared to the HEVC codec for similar perceived output quality. Given the uncertain future of adoption of VP9, there is little reason to suggest that 4K video data rate requirements will exceed the 25 Mbps rate, even with the lower compression efficiency performance of this codec compared to HEVC.

3.2.1 The 8K Prospect?

4K resolution is likely to make it into the mainstream computing paradigm quickly and dominate on all manner of devices. It is not difficult to imagine an 18cm (7") tablet device having such resolution, even with the current thought that such resolution might be 'overkill'. The question is perhaps not one of if, but when. Will it be 2 years, 5 years, or 10 years? We can make some guesses, but the safest projections of the future is for the prospects of 4K large-screen television units declining in price to the extent that there is no market for anything less than a 4K resolution TV screen above 102cm (40") size in as little as 3 years.

These rapidly declining prices are also likely to exacerbate the trend of consumers adding to the number of screens in a premises, rather than only buying new technology options as a replacement. The implication is that 4K technology penetration rates are likely to be significant. Consumers with 81cm (32") or 102cm (40") screens as their main device today might be enticed to purchase a 140cm (55") or larger 4K screen as the price descends through the \$2000 barrier. Couple this with the likely emergence of 4K content over the next few years and the prospects for 4K TV appear very bright indeed.

8K resolution at 7680x4320 pixels is also likely to establish a market presence, but it is far from obvious whether this will become a mainstream format. There are valid reasons to suggest that 8K resolution is overkill for the majority of consumers. On the other hand, we must be careful not to predict the future on the basis of outdated paradigms of the past. There is reason to expect that production efficiencies will extend to unit pricing for large 8K screens being low enough to tap a number of market niches. Manufacturers are unlikely to be sitting on their hands in R&D terms on 8K units, but most would no doubt be waiting to see exactly how the 4K market develops before placing major bets on 8K. They will most likely have the answers they need on the 4K market within the next 18 months to two years.

We can thus expect that a reconsideration of the prospects of 8K resolution in as little as two years' time will lead to relatively concrete conclusions. Today the crystal ball is decidedly cloudy on this question.

From a network development perspective the uncertainty of 8K TV is unfortunate, as this presents a large variable factor in our projection of future bandwidth demand. A compressed 8K stream is expected to require as much as 85 Mbps to transmit. This is close enough to the limits of a conceivable FTTN network paradigm to cause substantial concern. Widespread demand for multiple

4K streams per premises in a very small number of years can be predicted far more reliably. Potential mainstream emergence of 8K is a consideration, but changes our equation little if we already take the 4K multiple stream threat seriously.

3.2.2 3D Television?

Even with HD resolution, there is challenge associated with the extra bandwidth involved with 3D transmission. While 3D content is limited in extent (selected movie and sporting content only), there is little call for broadcast transmission. Globally, trial broadcast of 3D sporting events, including a dedicated Foxtel 3D sports channel have not resulted in the success hoped.

The economics of greater broadcast of 3D content are restricted by the practicalities of the restricted audience base. Many consumers have resisted the impulse to purchase a 3D capable unit, after having recently upgraded to a HD LCD display. For others the hassle of 3D glasses (especially for active 3D, but also for passive 3D) makes the experience unappealing.

The overall equation may tip slightly in favour of 3D for the larger screen sizes of 4K television. However, screens that are capable of the high refresh rates required for 3D and the high resolution of 4K only seem to be emerging on the market. On the positive side, passive 3D on a 4K screen compares favourably to active 3D in HD. We can expect to see 3D capable 4K TVs continue to make market inroads over the next year or two. The more immersive quality of a larger field of view coupled with 3D may be enticing in a number of market segments.

As suggested within the 2012 Analysys Mason report^[3], 3D television on a larger format screen where the screen field of view angle is significantly greater, brings with it additional implementation possibilities. A fertile area of research consideration is the provision of a more immersive 3D experience that provides variation in perspective as the viewer shifts their position either by head movement or full body movement. We can only speculate crudely as to the potential for such technology, but beyond the 10-year horizon we must allow that possibilities that today seem extreme may become a reality. The bandwidth required for such 3D implementations may indeed be very high, although the 300 Mbps suggested by the Analysys Mason report continues to appear outlandish. A more conservative approach is that significant processing capability at the receiver would be used to provide the extra 3D realism based on processing of image meta-data. We would be doing little more than taking a stab in the dark to attempt to specify the bandwidth requirement of any such future 3D 4K TV offering, but a bandwidth range of between 60 Mbps and 200 Mbps is not inconceivable.

Content recording in 4K 3D will require sophisticated equipment that is likely to be a step up in cost from already expensive 4K cameras, but this addition is unlikely to be a major contribution to the overall economics of the technology.

The 'threat' of 3D TV, and especially 4K 3D television is almost impossible to assess at this point. A clear pathway forward for 3D television may emerge over the next 3-5 years. Radio broadcast of 4K 3D TV content seems unlikely even if the broadcast industry adapts to the 4K paradigm. **From the present NBN perspective, if 4K 3D television was the only potential future demand we could safely suggest that it is far enough in the future to worry about at a later stage.** This suggestion is reflective of a paradigm of FTTN build today with mainstream upgrade to FTTdp or FTTP commencing after the end of the decade. **The reality is that 4K TV (non 3D) is here today, and this threat of demand to stream multiple 4K sources to a single premises is very real for the FTTN model for widespread deployment. Ultimately the conclusion is that**

the FTTdp option must be seriously evaluated if FTTP is indeed uneconomic.

3.3 The Evolution of Video Conferencing Technology

4K resolution screen technology is set to find wide application in video conferencing. All who have participated in large-scale video conferences will appreciate the resolution limitations of current technology solutions. High-end video conference solutions in dedicated equipment configurations are likely to be well supplied by NBN Co service offerings with sufficient reliable bandwidth provision. However it is the possible proliferation of generic use of video conferencing that is of most significant interest for NBN network provision considerations on the customer access end.

We do not have to imagine what it would be like to have Skype-type flexibility to the majority of desktops but with decent-quality video and audio feeds. It then does not take very much additional imagination to consider having a large number of simultaneous feeds from a distributed video conference group. With sufficient network capability, completely 'new' ways of working become possible. The much vaunted possibility of telecommuting may indeed start to become a reality for a wider section of the community.

A virtual office might consist of a large high-resolution (4K) screen with 16 or even more incoming feeds from desks of colleagues. Asking a question to a group or an individual can be as easy as yelling across the room or walking into someone's office. Such a scenario is not for everyone, but it provides real advantages and possibilities for many and enables mechanisms of competing on the world stage in a way that may not exist presently. We can not underestimate the potential benefit to the economy from enabling such new work practices. Freed from physical location bounds as never before, productivity may soar in a large number of niche areas.

Equally, a new paradigm can be expected to emerge for educational tutoring assistance from school to university levels. The seeds of these businesses have already been sown and are showing promise without the presence of highly capable network technology. It takes very little imagination to see how things might evolve over the next several years.

Parents currently ferry children far and wide after school to reach quality instrumental music teachers. The potential for a 'disruptive' introduction of technology creating wide benefit is significant.

The implications for bandwidth requirement of the above scenarios is highly variable. A high level of bandwidth is required for the most demanding possibilities. Significantly, we expect the above demands to emerge in an incredibly distributed physical nature throughout the nation. The conclusion is that whatever access technology path is chosen for the NBN, it needs to provide flexibility to meet a large demand variation from user to user. **Clearly FTTP has this capability while it is present in only a highly restricted sense with the FTTN model. FTTdp effectively matches the FTTP promise when we consider the reality of 250 Mbps VDSL2 speeds and the ease of upgrade to G.fast and individual FTTP on demand.**

One significant point to note is that the low scene activity levels commonly found in face-to-face video conferencing environments allows for a generally very high level of compression performance. However, the practical realities of the multi-user video collaboration scenario acts to limit high compression complexity. There are many variables to consider, and provision of effective large video collaboration systems receives research attention globally. A local example is provided by the Smart Services CRC, iSee Video Collaboration project^[20]. The future also promises demand for a high level of similar service capability in mobile device environments. Bandwidth limitations

on the mobile cellular network present significant research challenges. Some observers also appear to foresee widespread adoption of video collaboration technologies on a broader social basis within our society. We can not discount this possibility, but there is good reason to suggest that this outcome, at least, will not occur overnight.

3.4 Video Upload Requirements

With the growing popularity of cloud storage services and social and professional image and video sharing, it is worth contemplating the direction in which mass-market video capture is heading.

Cheap digital cameras available for purchase today generally have the capability of recording high definition video at high quality. The movement toward 4K display technology will only spur this competitive device market to continue to innovate with regard to supplying ever greater video capture capability. A similar trend is readily observable in terms of the capabilities of mobile device cameras moving rapidly from being a second-rate device to a superior image and video capture solution. Both mobile devices and newer digital cameras are Wi-Fi enabled to facilitate no-fuss upload of image content.

We can assume that the volume of content downloaded within any particular premises will generally exceed that uploaded by a significant margin, and that downloaded content will continue to be on average at higher bit rates than captured video content for the immediate future at least. Overall our conclusion must be that upload demand is likely to continue to experience rapid growth, but in the majority of cases this will lag download demand.

3.5 Small Business Upload Requirements

For many small business enterprises, significant upload capability is a major business enabler. For some small businesses, a high level of upload capability is needed at employee residences. Paltry upload rates on ADSL are insufficient for many. VDSL deployment in FTTN promises to greatly improve upload capabilities even for long copper loop lengths. However, for a significant number of small businesses a major improvement from a very low starting point is not enough. It is difficult to gauge how much upload bandwidth will be required by small businesses, but we can imagine that excessively demanding users will require a dedicated fibre connection. For more moderate usage we can imagine that 20 to 60 Mbps of upload capacity should suffice to keep most happy. Obviously this is just a hypothetical figure with a small amount of intuition thrown in. It is also important to note that upload bandwidth demands must be assumed to increase steadily over time.

Any technology option we consider for NBN deployment must be capable of supplying something similar to this level of upload capacity to a small but significant proportion of premises. Many small businesses are not anywhere near as concerned about high levels of CIR (Committed Information Rate) as they are about high best-effort bandwidth, and for some enterprises it is upload bandwidth that is key. In this sense the labelling of NBN Co plans as residential or business might be less than perfect. We can expect that the vast majority of smaller businesses will be more interested in best-effort bandwidth than CIR, especially considering what appears to be monopolistic pricing on basic “business class” plans.

4. The Contentious Issue of Contention

Determination of network utility from metrics such as Contention Ratio and Average Busy Hour Throughput (ABHT) is fraught with danger. The relationship between these metrics and network utility is highly non-linear in complicated ways and evolves over time. Engineers and decision makers alike can be easily misled by placing too much significance on such metrics.

Network link contention is a familiar topic for all. In our modern, mobile, and connected world we regularly experience network contention from a slow download on a hand-held device to glacial-pace web browsing. Those working in the telecommunications and networking industries without exception have a good general understanding of contention as it touches so many aspects in all areas of this broad industry.

The issue becomes 'contentious', however, when it is poorly understood and principles are applied incorrectly. The breadth and complexity of communications systems implies that extreme caution must be used when applying 'rule-of-thumb' methods of understanding contention. The network provisioning consequences are surprising in some situations, and can run counter to 'intuition'. It is particularly vital that key decision makers such as senior company management or strategic direction-setting teams, are made aware of some of the common pitfalls.

An introduction to the basics of contention is provided in Appendix A. This appendix includes a discussion of the statistical nature of user demand and concludes that the existing NBN FTTP architecture can be deemed to supply a 1 Gbps/400 Mbps level of service to all users, at least while actual demand for this level of connectivity is low enough for contention impacts to be minimal.

Contention issues become more significant in the satellite connectivity environment, and are a particular concern in the NBN fixed wireless scenario. These are discussed in the following sections. Contention has a clear interconnection with the concept of Quality of Service (QoS), and this is discussed in section 4.3.

4.1 NBN Co Satellites and Contention

The NBN Co satellite solution for remote premises provides for 700 Mbps down and 200 Mbps up within each spot beam^B. If we assume this serves an average of 2000 users, this presents a guaranteed minimum capacity of approximately 350 kbps down and 100 kbps up. During off-peak or non-busy times it is reasonable to expect that individual users can indeed experience a service close to a real 25/5 level. Should as little as 2% of the users in any particular area be attempting to use their full 25/5 capacity (attempting to download a very large file or stream 4K TV), then capacity limitations will clearly have an impact.

For remote properties there is likely to be little alternative to the satellite option, and without the provision of significantly increased amounts of spectrum, this is likely to be the best outcome possible. Clearly it represents a significant step up in connectivity for many premises over currently available options.

B This is the data rate shared by all users in the spot beam and can not be assumed to be available to a single user.

There is also the prospect that an enterprising organisation will supply a video broadcast caching solution tailored to the specifics of individual satellite beams to help alleviate the increasing demand from popular online video sources. However the economics of such a solution are perhaps marginal (technical requirements vary sufficiently from mainstream IP TV product and service offerings), and it would most likely require regulatory cooperation of a type unfamiliar to the industry. Where the nation has novel problems we should attempt to support the creation of novel solutions. Fortunately in this case the full emergence of this 'problem' is still a couple of years away. It is perhaps timely for the Department of Communications to begin considerations, but is outside NBN Co's focus at present.

With dominant usage paradigms centred around web browsing, email, and facebook, usage demands are relatively bursty in nature. Over a large catchment of 2000 subscribers this bursty individual behaviour translates to a more-or-less constant bandwidth demand with temporal evolution over the course of a day. In the majority of cases users perceive the effect of heavy use only by way of slower loading of web page elements or slower download progress. It is when users attempt to stream high-rate content such as video that they appreciate the full limits of the network. Crucially, even standard definition compressed video streams are likely to cause difficulty with the satellite network at a time when embedded HD video content is becoming prevalent on the web and Netflix is leading the way to 4K content.

4.2 NBN Co Fixed Wireless and Contention

In recent years we have seen a steady stream of ill-informed media comment about the future promise of wireless technologies. The reality is that limitations of wireless technologies are likely to continue to be a significant concern for both mobile networks and Wi-Fi in the premises^C.

There are some more sparsely populated areas where fixed wireless might make sense as an access network option, perhaps accounting for a few percent of the nation's premises. However careful analysis is needed to identify such areas. The 2010 NBN Implementation Study appears to have recognised this as a concern and noted that areas that were already well-served by DSL technologies would be unlikely to be better served by fixed wireless. The reality is that the shared link nature of wireless service is seen to be a severe limitation based on average usage assumptions of today. Once we allow for growth in average demand or the existence of a small number of above-average users on a single wireless sector, the limitations of the fixed wireless option become more extreme^D.

Where there are no alternatives to fixed wireless, then, as with satellite, this is perhaps the best outcome we can hope for. However, in many regions DSL provides a very real alternative to fixed wireless. Those regions with dense premises clusters around townships are particularly well suited

C In crowded Wi-Fi environments such as high-density urban living or working settings, the last Wi-Fi link will perhaps be the site of more contention than any other part of the network if data traffic demands continue to increase as we expect. The solution to this problem of Wi-Fi congestion is relatively simple but may involve a mindset shift on the part of the general user base. For many years Wi-Fi has been seen as a trivial solution to provide network connectivity within homes and businesses. It will continue to provide this capability for a large number of users, especially when we consider the wider bandwidth capability of Wi-Fi using the 5 GHz spectrum. However, a growing number of users will need to consider other network media such as Cat 6 cable or ethernet over powerline.

D It is possible to improve fixed wireless outcomes through the use of large increases in RF spectrum. This option may need to be pursued as demand grows for those areas where fixed wireless is the only viable access network choice. For areas where VDSL deployment provides a better outcome, VDSL is still likely to provide a better outcome even with large spectrum allocation to a fixed wireless option.

to DSL implementations. Assuming the existence of fibre back-haul to a centrally located exchange, then many regional towns can expect to obtain a significant boost to connectivity levels by VDSL2 upgrades. The lack of contention on the copper loops translates to a vast improvement over fixed wireless for the majority of users, especially those very close to the exchange.

For a more distant premises, say at 2.5 kms from the exchange, we might conservatively assume a 10 Mbps or lower download rate over VDSL. The question as to whether this 10 Mbps rate (with no contention) provides greater capability to an individual user than a 25 Mbps contended service level over fixed wireless is a key consideration. To be able to answer such a question authoritatively we need to identify the average usage characteristics for the rest of the users in the wireless sector, and understand the usage demands of the particular user. Under full contention a wireless sector with 60 users provides only 660 kbps down and 150 kbps up. Clearly the VDSL 10 Mbps provides significant advantage compared to the full contention case, but this must be balanced against the benefit of access to the higher 25 Mbps rate available through fixed wireless when there is limited impact from contention.

To an individual user the network utility afforded by a guaranteed 10 Mbps connection is significant compared to the prospect of 660 kbps. The additional utility afforded by the prospect of a 25 Mbps connection is not negligible, but is most certainly smaller than the jump to 10 Mbps. When all these factors are considered in conjunction with real premises distribution patterns, it is clear that many regional towns will be immensely better served both now and in the future by a DSL-based solution compared to the fixed wireless option.

Of significance for such an assessment is the reality of video streaming services being provided by set-top boxes over ADSL today. Such services are being actively marketed to regional areas where NBN Co is currently deploying fixed wireless. The same level of service provided by DSL to these set-top devices simply can not be supplied by fixed wireless under assumptions of even moderate consumer uptake.

VDSL for small regional towns provides the flexibility for regional business growth. Many such communities have creative cottage industries. Some of these industries can expect to see significant future bandwidth demand as they evolve with the emergent modern global economy. A reliable 25/5 Mbps service will facilitate many such regional businesses to participate in the evolving global economy, but limiting such areas to contended 25/5 service via fixed wireless prevents local growth initiatives as contention levels increase with demand evolution and growth. Provision of high-speed networking via VDSL2 in areas close to the exchange opens the door to future community-funded network extension via FTTN, FTTdp, or even FTTP. Important options for the future are provided to such communities in a way that fixed wireless can not deliver.

One of the 'advantages' of fixed wireless that can not be overlooked is that it so severely limits bandwidth demand by contention on the local RF sector, that the provision of 'back-haul' becomes considerably simplified. Indeed, the NBN Co Network Design Rules^[5] suggest that a single tower of three sectors of 40 Mbps nominal capacity (60 Mbps for premises with high SNR able to employ high-rate modulation at 64 QAM) can be served by a 100 Mbps microwave link. Naturally a VDSL2 deployment unleashes demand to the extent that such a microwave link is likely to be completely inadequate. Higher capacity microwave links are a possibility, although without high towers the attractiveness of this option is limited. Fibre connectivity to the exchange is thus important.

It is unclear how many such regional exchanges currently operate ADSL over a bundled copper back-haul link. In these cases, unleashing latent user demand by allowing higher speed premises

connectivity through VDSL2 will quickly result in greater contention problems on the limited back-haul link. The additional cost of upgrading the regional exchange connection is still likely to be favourable to the VDSL option compared to the high cost of fixed wireless tower deployment. In some specific cases we can expect that a more thorough analysis will be required in order to determine the optimal deployment solution.

We can safely assume there are several hundred regional locations currently planned for fixed wireless implementation where VDSL options need to be properly considered. Using the right metrics and tools it is conceivable that a single small team can individually assess each and every one of these locations and make a sensible determination as to the best access technology to deploy. **The limitations of fixed wireless as an access technology are very real and become ever more so with growth in traffic demand. In building the network of the future we must avoid blanket categorisation of areas. The outcome of investment in the NBN can be maximised by careful deployment selection.**

In areas with significant amounts of foliage there are additional concerns related to the fact that 2.3 GHz radiation does not propagate well through trees. The existing NBN Co paradigm of defaulting users with no 2.3 GHz fixed wireless to satellite adds to the satellite loading problem. Where good VDSL capability potential exists it is insensible to forgo this option in favour of adding traffic to the satellite.

Where NBN Co Fixed Wireless infrastructure has been deployed it will no doubt be used by some customers. However, the comparative benefits of the copper pair network may encourage commercial VDSL deployment as a competitive offering. There are revenue risks to NBN Co that are likely to be most severe in the regions where fixed wireless is overall a poor technology choice. In such areas the price equation may be sufficient to ensure good competitive VDSL deployment as global VDSL uptake continues and vendor equipment pricing declines. However, monopolistic effects also have the potential to limit such competitive service offerings and the medium term outlook is for NBN Co fixed wireless to provide limited service competition in such marginal technology deployment areas. It is thus difficult to predict with any clarity, but with very low subscription to NBN Co fixed wireless service in a particular location, the operation costs alone would dictate service decommissioning consideration.

Fixed wireless has a role to play in the NBN, but it must be deployed with careful consideration to suitability for any given area.

We can not honestly convince ourselves that data usage statistics applicable a decade ago apply with modern network demand. Video services supplied over ADSL are being vigorously marketed to regional areas today. The step change in data usage and network expectation that such services imply is completely incompatible with the NBN Co fixed wireless option. Fixed wireless very clearly has limited appeal where other delivery options are viable. **We must take this threat seriously and ensure the nation doesn't throw good money after bad with continued fixed wireless deployment in inappropriate locations.**

It is worth noting that bonded LTE and VDSL product offerings are in existence. These provide for a high level of peak throughput via the high peak rates of LTE, combined with the dedicated capacity of VDSL. One way this product offering can be envisaged is the flexibility of LTE coupled with a high CIR (Committed Information Rate) that can be supplied by VDSL. On the surface it appears that such an approach might have long-term value in NBN Co fixed wireless areas, especially those areas where poor deployment commitments have potentially already been made and severely under-utilised NBN Co fixed wireless assets are likely to exist in the near future.

However, the reality is that without the benefit of high levels of MIMO (Multiple-Input Multiple-Output), LTE just doesn't offer a significant enough level of service to make a bonded LTE/VDSL product exciting. The lack of significant MIMO advantage is one reason why regional fixed wireless doesn't live up to the expectations that seem to have been implicit in early NBN Co considerations. More multipath exists to facilitate MIMO in highly built environments, hence the NBN Co fixed wireless LTE limitation to a 60 Mbps channel (40 Mbps for those slightly further away and using 16-QAM rather than 64-QAM) and not the headline LTE rates of 150 Mbps or 300 Mbps.

Having said this, a bonded LTE/DSL product may provide a much-needed boost to those customers distant enough from an exchange-hosted VDSL deployment to have a low bandwidth outcome. If we assume the large majority of premises within the region are densely clustered around the central exchange (as they are in many towns), then these premises will be well-served by VDSL. The remaining small number of customers are likely to obtain value from the shared fixed wireless service. If subscriber numbers are low enough for fixed wireless (a handful per sector), then even with a considerable increase in usage demand profiles (as is destined to occur) a worthwhile service level can be obtained from fixed wireless, with VDSL providing an important improvement to base-level throughput. We can not escape the conclusion that fixed wireless is a poor technology deployment choice for areas where the majority of premises are clustered tightly around a node/exchange, but if the infrastructure exists it is worthwhile determining what level of use can be obtained prior to considering writing off the asset and decommissioning the facilities.

NBN Co fixed wireless may indeed be the only viable option for a significant percentage of premises in Australia. However, where VDSL implementation is viable it must be pursued. A proportion of regional areas have dense township premises clusters ideally suited to VDSL. In these cases VDSL service promises far greater overall network capability for potentially less cost than fixed wireless. VDSL also provides network infrastructure that can be readily expanded to meet future growth demands.

4.3 QoS (Quality of Service)

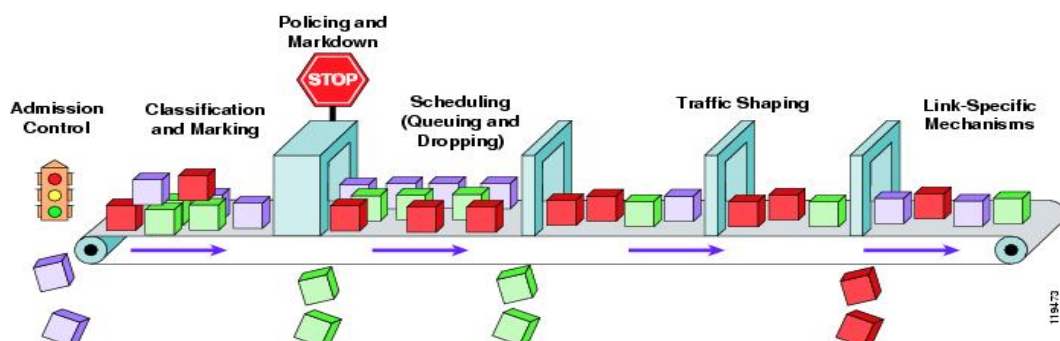


Figure 4: QoS implementation involves the 'whole' network. (Image Credit: Cisco)

QoS typically refers to the provision of different traffic classes to ensure application data demands are optimally met by the network across a variety of scenarios. Some applications such as voice and video telephony (teleconferencing) are inherently two-way and have limited ability to tolerate packet delay. Contention on any one network link is often handled by buffering at the network node, and there is the ability to prioritise traffic in the outgoing buffer according to its designated traffic class. Multicast capability is sometimes considered as an additional traffic class and it has implications for node buffer implementations.

In the GPON example, link contention is likely to be felt first at the 40 Gbps link to the transit network. Traffic heading in both directions can be prioritised in the buffer according to traffic class at the relevant sending end of this 40 Gbps link. There is also the chance that a contention limitation occurs on the GPON fibre, although the contention ratio is lower in this case.

The major network implementation concern with QoS implementation is not the provision of prioritised buffering as such, but ensuring this is done in a managed and effective way.

Naturally provision of end-user ability to mark packets as high priority without limit leaves a system that is open to abuse. The network must manage and control QoS. One of the obvious ways to do this is the approach taken by NBN Co in offering products based upon different levels of committed bandwidth (CIR – Committed Information Rate).

Base-level (domestic) NBN plans incorporate a level of CIR apparently primarily designed to facilitate telephony services. 150 kbps CIR presumably allows two simultaneous G.711 PCM conversations at a raw data rate of 64 kbps plus packetisation and signalling overhead. Such plans are readily transferable to the FTTN or FTTdp implementation environments.

In the FTTdp scenario the contention picture closely matches that of GPON, and traffic class management can be expected to follow the FTTP situation. On the downstream side any traffic class prioritisation for data heading to the FTTdp nodes occurs at the same location as the centralised active FAN equipment for the FTTP case. On the upstream side the addition of traffic class prioritisation leads to trivial implementation requirements for a 150 kbps CIR. Higher levels of CIR present different issues as mentioned below.

In the FTTN case, traffic class management on the back-haul link can follow the FTTP example. As separate channels are used to each premises from the node over the copper pairs, there is no contention on this side of the network. We must note that crosstalk does cause the effective link rates to vary. For as little as 150 kbps CIR this variation is unlikely to be of concern except in extreme pathological cases (where low VDSL throughput is achieved presumably due to distance or copper pair quality issues, very high crosstalk exists, and vectoring fails to provide any significant gain).

In both the FTTdp and FTTN cases, higher levels of CIR to provide additional traffic classes or additional priority one traffic, can cause implementation concerns. In the FTTdp case the short copper loops ensure that enough link bandwidth is present (combined with limited crosstalk implications in the majority of cases) to enable delivery of significantly higher levels of CIR. The implementation concern with FTTdp is just ensuring that vendor equipment has sufficient buffering capability to provide the desired QoS solutions. Compared to GPON the buffering implementation is significantly less obvious due to the variable nature of the link rates involved, and there is less economy of scale with the buffer location at the micro node device. (Buffers at network points where there is higher fan-in can benefit significantly from the effects of statistical averaging.)

In reality, for FTTdp these matters are a mere trifle. The high link throughput related to short copper length provide substantial capacity for QoS implementation. Should there be a situation where a

particular customer requires a level of QoS capacity not able to be supplied by FTTdp, the simple on-demand upgrade path (at relatively low end-user cost) to full FTTP provides flexibility to cater for all demands.

The lower levels of link bandwidth in the FTTN case provides some concern with regard to being able to supply higher levels of CIR to all customers. Qualification of customers for higher CIR based plans is a network management problem involving tweaks to planned business and operational support software (OSS/BSS). While these problems are not insurmountable, they are clearly an inconvenience and must be carefully assessed as to their cost implications.

The other very real question with FTTN implementations is whether the node equipment supplied by major vendors is QoS ready. Crucially, NBN Co requires QoS support of a relatively specific nature as NBN Co business service plans are a significant part of the company's projected income stream. It is far from clear whether global FTTN deployments have enough synergy of interest in particular QoS implementation to drive equipment vendors in the direction that NBN Co would like. This presents an important risk item to NBN Co. As FTTN has been considered in some detail internally as a part of the current on-going strategic review process, it is expected that these matters are being considered carefully within NBN Co. With luck NBN Co already has sufficient answers to questions of QoS provision for FTTN.

While provision of a CIR for voice services (telephony) appears necessary, there has been persistent questioning of the need for broader QoS from some sectors of the industry. The large monthly charges NBN Co associates with even moderate levels of CIR helps fuel some of these questions. Modern fibre-based communications networks are able to be provisioned to have minimal contention even under most peak load conditions. Marginal costs of providing slightly higher levels of CIR to business users is not large. The nation needs to have fair and equitable pricing schemes that encourage economic growth. Our nation has suffered greatly in terms of communications infrastructure provision and flow-on economic benefits in the past from a strong telecommunications industry focus on using monopoly power to extract premium prices.

While NBN Co must create an income stream to help fund network deployment costs, extreme care must be taken to ensure that pricing does not artificially inhibit service adoption or suppress demand. Particular consideration should be paid to those areas where high levels of services can be supplied cheaply, generating strong revenues (with sensible pricing structures) to subsidise higher cost areas.

FTTdp allows substantial flexibility to supply high levels of service (high CIR) to drive revenue growth (with appropriate pricing). FTTP offers the same or higher capability (at higher cost).

5. NBN HFC Incorporation

There is little doubt that existing network capabilities must be fully exploited to maximise the level of national network utility provided during the prolonged NBN build period. There are strong arguments to be made that HFC networks are able to meet most current network demands in their footprint areas and are certainly able to meet foreseeable demand over the next couple of years with presumably minimal infrastructure investment. With more moderate levels of investment there is every reason to suggest demands can be well supplied by HFC for the majority of customers over at least the next decade.

There is good reason to expect that HFC networks can readily supply wholesale open access layer 2 service similar to NBN Co product offerings to date. As such, NBN Co is readily able to absorb HFC infrastructure into its product mix, and focus can turn to areas where maximum network utility improvements can be provided to the most people in need within the smallest cost/time budget.

The 2010 NBN Implementation Study^[2] suggested that existing HFC networks needed careful consideration with regard to being incorporated into the NBN product mix. The desire for a 'pure' FTTP NBN is perhaps one of the primary reasons why HFC has effectively been dismissed until the December NBN Co Strategic Review report^[1]. HFC networks can provide high levels of service capability and can not be summarily dismissed. The recent release of DOCSIS 3.1 specifications from CableLabs^[7], provides for a maximum of 10 Gbps downstream data channel and up to 2 Gbps upstream (with greater bandwidth allocation the capabilities are even higher). Cable-based access technology clearly must be given a fair hearing. The published Strategic Review document suggests that HFC has a substantial role to play in the NBN, but what is missing is a clear picture of the long-term evolution strategy that applies for HFC deployment areas.

It is important to note that some current users of HFC networks report relatively poor connectivity outcomes.^E It is not difficult to find speedtest.net results quoted in various web forums of the order of 18 Mbps down and 0.4 Mbps up. While on the downlink side this is comparable to ADSL2+, the relative reduction in uplink is significant. Modern domestic user demands employing significant upload of photos and movies to cloud storage are severely limited even at the best ADSL2+ upload rates. Reports from early users on 25/5 Mbps plans in fixed wireless areas highlight upload improvements, reinforcing the conclusion of many industry commentators over the last decade that low upload speeds are a major issue. The upload speed issue is especially important for many small businesses.

Speculation abounds that part of the reason why HFC network capabilities are reported, by some, to be lagging demand is the prospect of network closure having prevented even minor capital investment activities in recent times.

E Presumably these reports are from users primarily attempting to use the network during peak times as other reports are positive on network capability provided by HFC. An additional complication could be an incomplete process of transition to DOCSIS 3.0.

Winning the hearts and minds of HFC users and others in HFC deployment areas will be important for NBN Co. A detailed and viable technology road-map must be established to produce superior cable connectivity outcomes.

Fortunately there are also many positive reports of solid connectivity results being achieved in cable areas. HFC users are often very happy to report superior results compared to ADSL connectivity. Properly operated, managed, and maintained, HFC networks have ample capacity to meet user demand that exists today and over the next several years.

The long-term migration plan for HFC networks is naturally a crucial technology matter to have mapped out to ensure that any additional investment in HFC networks to improve capability in the short term translates to long-term value where possible. Compared to GPON or even FTTP approaches, HFC use represents some amount of engineering compromise in dealing with legacy issues in what was originally envisaged as a broadcast-only medium. Recent DOCSIS releases and other developments such as Huawei's cable G.fast announcements make it clear that HFC has real capabilities.

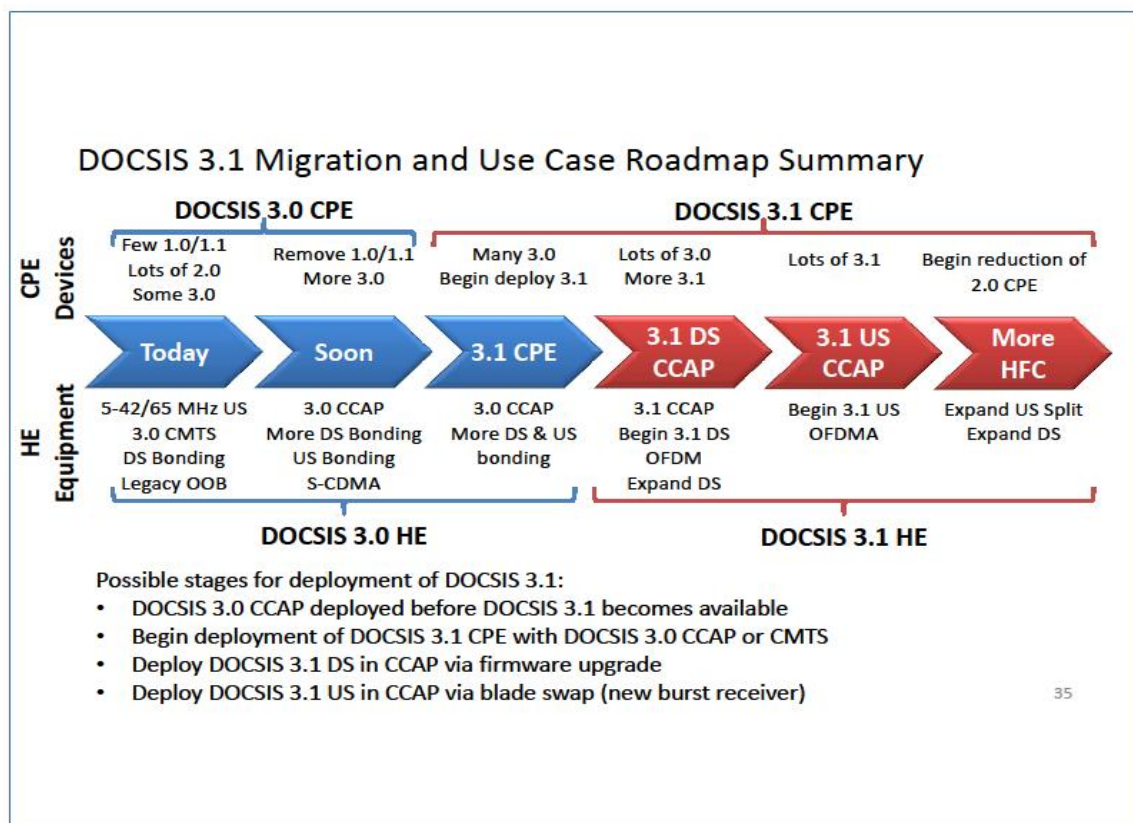


Figure 5: DOCSIS Migration (from [21])

Figure 5 presents a graphical picture of the upgrade road-map from DOCSIS 3.0 to 3.1 (extracted from a webinar hosted by CED and sponsored by ARRIS and Cisco^[21]). Here the most crucial point is that interoperation of 3.0 and 3.1 equipment allows a flexible and staged upgrade process. Note also that DOCSIS 2.0 CPE can continue to operate on the network with DOCSIS 3.1. The slide also indicates that deploying the downstream side of DOCSIS 3.1 in the CCAP (Converged Cable Access Platform) is a process of a firmware upgrade, and as such is a relatively painless process (quick and low cost).

The predominant vision in the cable industry is a gradual movement toward full DOCSIS 3.1 deployment over a period of some years. Importantly, DOCSIS 3.1 technology has been designed to interoperate fully with DOCSIS 3.0. DOCSIS 3.1 CPE can be deployed immediately upon devices becoming available from manufacturers, independent of any upgrade to 3.1 at the CCAP Head End. After 3.1 CCAP upgrade, 3.1 service can be gradually introduced as more CPE supporting 3.1 is connected.

DOCSIS 3.0 and 3.1 provide what are considered today to be impressive levels of connectivity. However the shared nature of the channel medium implies that as user bandwidth demand grows and demand statistics evolve, contention issues will continue to present a problem. HFC node splitting is the obvious way to resolve such concerns on a short-term basis, but it is far from obvious that node splitting alone will help migrate HFC to what might be the 'end game'.

For a longer-term outlook there is significant interest in the concept of FTTP being the ultimate pinnacle of evolution for extant networks. The EPON standard provides the basis for this FTTP vision, and it is widely recognised that legacy video streaming approaches (the bread and butter of cable companies to date) is giving way to flexibly managed IP video content. Globally, where cable companies are looking to extend footprints, fibre deployment and RFoG (RF over Glass) is seen to be highly competitive against the option of installing new coax cable. The vision of HFC is then one of a global industry ecosystem driving toward high broadband data capacities with all the 'usual' requirements of low latency, traffic management, flexible upstream versus downstream profile, economical installation options, etc..

We can expect HFC networks to be responsive to market demands in continuous improvement of network capabilities. This assumes an appropriate level of HFC-specific technical and management oversight within NBN Co.

5.1 The Immediate Future and Beyond

At this point it seems prudent to focus on immediate requirements in the HFC network areas such as making a layer 2 wholesale product available, upgrading to DOCSIS 3.0 and 3.1 as appropriate, performing node splitting where needed, and filling in footprint black spots. As this work proceeds over the next 2 to 3 years there is ample time to formulate a detailed plan for long term network evolution in HFC areas. Such detailed planning must commence immediately and be performed with adequate input from domain specialists.

Using the time over the next 2 or 3 years to implement low-cost, 'obvious', improvements in HFC footprint areas allows careful consideration, with expert input, of the longer-term HFC evolution

question. Opportunities to ultimately evolve to a common technology platform over the entire NBN may seem remote, but at the very least we must seek the opportunities to provide communications product services that are largely indistinguishable. We must also be open to the possibility that HFC areas might eventually be overbuilt by FTTdp or FTTP deployments. In such a scenario any additional investment today in HFC assets has a real risk of becoming stranded. The HFC technology road-map is vital.

Many of the concerns of HFC, such as the provision of sufficient guaranteed throughput (CIR) for Traffic Class 1 services (or TC2 and TC3 for other NBN Co product offerings), are not new and have answers available from experts in the cable industry. A large list of considerations can already be assembled by reference to comments contained in the 2010 NBN Implementation Study Report. Recognising that real cable engineering expertise is difficult to find (at the level required) in the Australian telecommunications industry context, NBN Co must actively seek qualified external input. Careful attention is required to such a task, as there are many suggestions that NBN Co has failed to recruit suitable detailed technical expertise in the past.

The legacy of HFC is undeniably sophisticated/complicated compared to a new FTTP or FTTdp deployment. Detailed domain knowledge is required to ensure that any non-standard industry approaches are the correct ones. The equation is further complicated by the reality that there are two quite different legacy HFC networks in existence in Australia. These networks operate variously on the 8 MHz channel bandwidth derived from Europe or the 6 MHz channel bandwidth derived from North America, and employ markedly different physical layout topologies. DOCSIS 3.1 upgrade promises to ameliorate the channel bandwidth differences, potentially allowing some form of interoperation between the two networks. Crucially, DOCSIS 3.1 CPE is identical, facilitating migration of users between networks should that become desirable.

To maximise value to the nation in terms of full exploitation of available resources while ensuring a winning long-term outcome, it is entirely conceivable that novel approaches may be required on some issues relevant to HFC network deployments. With the culture of Telstra Research Laboratories endeavour now largely extinct in the local telecommunications industry, such an R&D challenge presents as a major undertaking. It may however be an undertaking we must engage in to maximise future network capability without undue cost implications. NBN Co, with an effective and functional senior management team, has a good chance of being able to provide sufficient focus to such issues with recruitment of appropriate levels of external expertise.

While it has been widely stated that the initial intention was for the Optus HFC network to be decommissioned once customers had been migrated to the NBN, there has not been a clear vision articulated for the long-term Foxtel plans. We can expect to see Foxtel continue to move to a service model as opposed to a vertically integrated content supplier reliant upon their own network infrastructure. Under the previous NBN plan of FTTP, there is strong reason to suggest that over time existing HFC connected Foxtel customers would migrate to services supplied over the FTTP connection. At some future point a commercial decision by Foxtel would most likely determine to decommission the HFC asset. Where HFC is used to supply NBN connectivity there is presumably less incentive for Foxtel to migrate customers from the legacy services. This goal may require NBN Co attention from the perspective of streamlining processes by only having a single model of customer access.

With up to one third of the NBN expected to be provided by HFC in at least the short term, we must take provision of a technology road-map to be a very high priority. We need to resist the temptation to delay such a crucial detailed engineering and business case investigation.

5.2 Some of the Many Questions for HFC

A brief list of points needing careful consideration in terms of HFC follows. This list should clearly not be assumed to be exhaustive.

- How is an open access layer 2 NBN product offering supplied in HFC areas, and what timetable is applicable for delivery of such NBN Co wholesale products?
- How should the two competing HFC networks best be exploited? Is there any prospect to merge the architecturally disparate networks?
- How is the throughput/latency/packet size issue best managed as network demands continue to increase over the coming decade? Does reduction of propagation delay by shortening coax cable lengths translate to any significant benefit?
- Can sufficient guaranteed throughput (CIR for TC1) be provided for immediate widespread deployment of wired telephony services to allow for decommissioning of the local copper PSTN? (Presumably the answer is yes, but this requires detailed confirmation.)
- Can guaranteed throughput be increased quickly enough to satisfy other application and NBN Co service demands?
- Is the Unsolicited Grant Service (UGS) approach scalable beyond the demands of voice telephony?
- What level of point-to-point business class services can be provided? Are network-based encryption schemes adequate, or will additional end-to-end encryption be required to avoid the physically simpler task of eavesdropping in the cable environment?
- Are DOCSIS QoS traffic class management techniques suitable for direct use by NBN Co product offerings, and will the overall QoS experience be similar to NBN Co product offerings on other network topologies?

5.3 HFC Black-spot Infill

HFC networks are reported to cover large areas of the major capital cities and areas such as the Gold Coast. However, there are some “HFC black-spot” areas within the confines of the larger HFC envelope. The NBN Co Strategic Review document suggests that new HFC infrastructure should be supplied to provide coverage to these areas. While this is an option that merits careful consideration, other options such as FTTdp deployment in these regions must also be investigated.

Without detailed knowledge of existing HFC network coverage it is difficult to characterise common reasons for black-spot areas. We can only speculate that the economics of deploying

infrastructure to particular areas did not warrant network extension during the initial build or network augmentation phase. While general levels of demand for broadband services has increased dramatically in recent years, it is possible that the nature of such areas that created an uneconomic situation initially, remains. Such considerations might be crucial to ensure sensible prioritisation of NBN Co activity over the extended network build phase.

Deployment of new overhead or underground HFC infrastructure is not likely to be materially different in cost and time compared to provision of FTTdp. With the later option there is a significant simplification in the premises connection as this can be implemented from the lead-in pit. The HFC option also leaves the question of legacy analogue telephony migration. The high level of network utility offered by FTTdp can be matched for average users by DOCSIS 3.1 in the short (or perhaps medium) term, but as demand statistics evolve the limitation of the shared channel nature of HFC implies a systematic process of node splitting to limit contention. We must then ask whether a deep fibre deployment in the form of FTTdp is a sensible alternative.

The recommendation is that the options of FTTdp and HFC must be weighed carefully for black-spot infill areas.

For very small gaps in HFC coverage (hardly describable in terms of a black spot), running new HFC infrastructure is likely to be the only sensible option. We can expect a sufficiently detailed analysis to be definitive on such matters.

6. Preliminary FTTdp Discussion

We can draw a strong conclusion that **FTTdp (Fibre-to-the-Distribution-Point, with this in general defined as the lead-in pit in the street outside the property boundary) must be considered in detail as a major technology option for NBN deployment.** It is entirely reasonable to suggest on the basis of a simplistic initial analysis, that a FTTdp model should be used for a large portion or the majority of the premises to be covered by the NBN.

FTTdp in the sense of deploying fibre to the street similar to the LNDN for FTTP, offers a high level of connectivity via VDSL2 over short copper loops, and offers flexible capability to meet variation in premises demand and future demand increases. A preliminary (back-of-the-envelope) cost analysis suggests that it may be cost-competitive with FTTN deployment. FTTdp is well suited to upgrade via G.fast deployment and individual FTTP on demand.

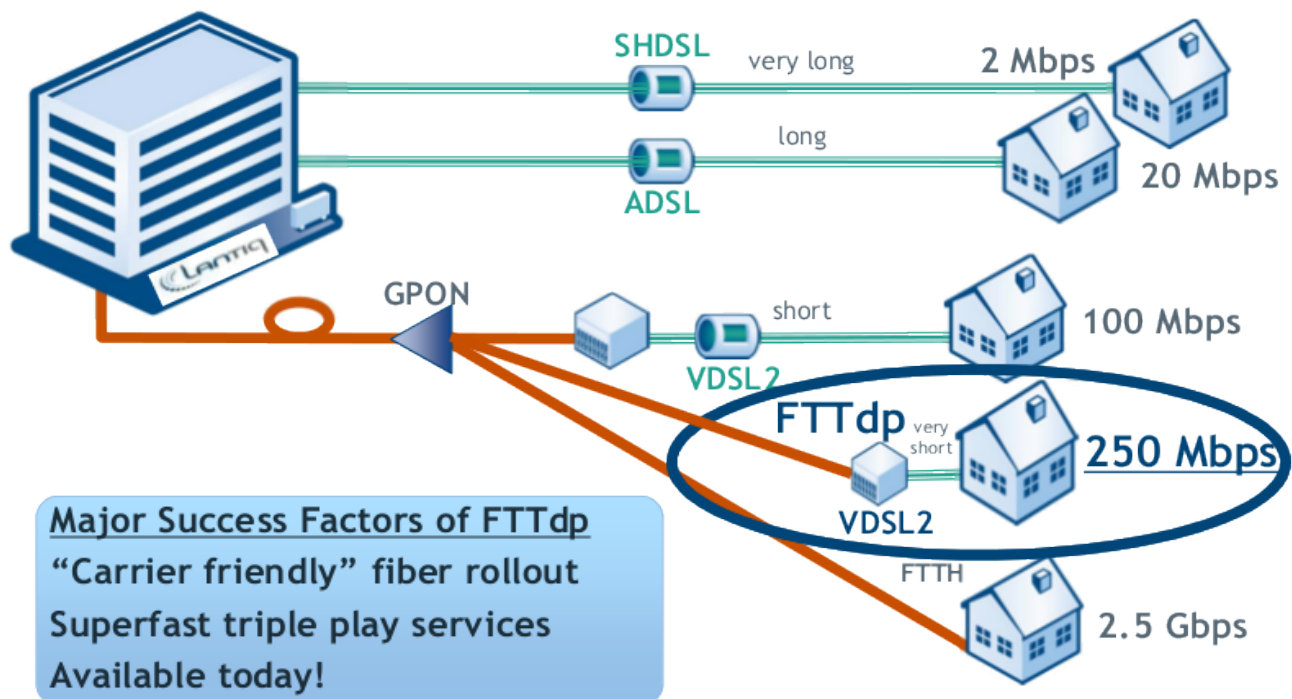


Figure 6: The very short copper scenario of FTTdp (from [18], Lantiq)

Figure 6 above displays the FTTdp scenario where fibre is deployed to the street lead-in pit, leaving a very short distance for the copper lead-in. Over such short distances VDSL2 can obtain speeds of up to 250 Mbps, using techniques of bonding and phantom mode. This diagram was extracted from [18], a Lantiq presentation on FTTdp. Lead-in copper lengths can be as short as an average length of 30m, making them ideal for not just VDSL2, but potential future G.fast deployment.

6.1 The G.fast Promise

G.fast^[15,16,17] employs several sophisticated techniques to extract the maximum value from legacy twisted pair copper networks. Initially it uses a much larger analogue bandwidth – 106 MHz (with

future versions at 212 MHz and 424 MHz) compared to the 17 MHz or 30 MHz bandwidth of VDSL2 profiles. It incorporates the proven VDSL techniques of bonding, vectoring, and phantom mode to maximise the throughput bitrate, and is widely reported to achieve 500 Mbps each way over short copper loops. As duplexing is done in the time domain, G.fast is capable of symmetrical upstream and downstream usage, but is crucially adaptable to varying data direction demand. TDD (Time Domain Duplexing) also ensures that there is ready ability to support low power states when equipment is not in active use. Rapid adjustment to channel characteristics in each frequency band, in a process labelled Transmitter Controlled Adaptive Modulation (TCAM), is noted to gain as much as 100 Mbps extra throughput over the entire bandwidth. **G.fast bridges the gap between VDSL and fibre connectivity, supplying sufficient bandwidth to premises for all but the most severe predictions of future demand (over very short copper loops).**




Figure 7: The G.fast Bandwidth Advantage

Figure 7 above (extracted from [6]) illustrates the advantage of greater bandwidth usage that is exploited by G.fast. As the frequency increases there is a reduction in the spectral efficiency parameter, but the channel data rate is proportional to the area under the curve and here we see a substantial gain compared to the VDSL2 baseline 17 MHz profile.

G.fast is anticipated to be ratified by the ITU-T in the early part of 2014. Early adopter deployment will soon follow. However, the interoperable nature of G.fast with VDSL2 (and legacy ADSL) implies that deployment can already begin in anticipation of future G.fast deployment for bandwidth-demanding customers as traffic growth occurs. The G.fast bandwidth (spectrum) allocation can commence above a 17 MHz or 30 MHz floor, ensuring full interoperability with VDSL2 deployments. Should all users on a micro-node migrate to G.fast from VDSL2, the G.fast bandwidth can subsume the lower vacated VDSL allocation without risk of crosstalk implications for other VDSL use in nearby areas.

A FTTdp deployment using micro-node devices implementing VDSL2 can thus commence today using proven equipment from a number of potential suppliers. G.fast support can be added with the addition of parallel or replacement micro-node technology as demanded. Should mainstream

demand for G.fast develop over the next few years as the network deployment is continuing, substitution of micro-node devices being deployed in new build areas with off-the-shelf support for G.fast may be a future decision.

The upgrade path to G.fast is therefore direct and relatively inexpensive. The same is true for individual supply of fibre drops on an as-demanded and user-pays basis. A FTTdp deployment is thus technically feasible today. There is no need to wait for mainstream G.fast deployment in other countries before commencing FTTdp even if we assume G.fast is likely to have a major presence in the NBN of the future.

Real-world performance of G.fast is yet to be proven in large scale. A concern is the substantial crosstalk problem related to the higher analogue frequencies. Even with sophisticated vectoring (CPU and power hungry), it remains to be seen how well-behaved G.fast can be in situations where there is a significant amount of crosstalk interference. For the FTTdp application, the relative absence of shared lead-in presents the best chance of maximal G.fast outcomes.

G.fast shows comparatively little potential for FTTN implementation due to the longer copper loop lengths and the larger crosstalk influences. We also need to note that the large gains that stem from bonding and phantom mode, applicable to both VDSL2 and G.fast, are unlikely to provide maximal benefit in the FTTN setting. The reasoning here is while spare pairs often exist in premises lead-in cables, we usually see a shortage of pairs from nodes to street pits, depending on geographical area growth profiles and the level of historical infill development. Even where spare pair capacity exists between the node and the pits, exploitation of this capacity is a potentially laborious and hence costly exercise.

We can safely conclude that G.fast provides a FTTdp deployment with substantial flexibility to meet the demand posed by unknown traffic growth scenarios. Initial FTTdp deployment is likely to remain VDSL2 based. Bonding, vectoring, and phantom mode already provide an ability for VDSL2 to obtain significant data rates over shorter copper lengths (up to 250 Mbps).

6.2 The Cost of FTTdp Deployment

Part of the reason why FTTdp and the promise of G.fast has received substantial global industry interest is that it shows some promise of being a cheaper deployment option than FTTN. If this proves to be true, then FTTdp is the obvious choice for implementation. Substantial network capability is provided by VDSL2-based FTTdp along with simple upgrade paths to G.fast and FTTP on an as-demanded basis.

Naturally FTTdp involves a significantly deeper deployment of fibre into the access network than does FTTN. However the cost of this fibre deployment is balanced against the higher CAPEX of copper remediation, time-consuming copper patching installation work at the node, and tricky deployment cut-over practices. We perhaps do not even need to take the OPEX differences into account for working with a greater proportion of the legacy copper network for the decision to become clear.

The NBN Co Strategic Review document indicates a customer connect (drop) cost per premises for

FTTP of \$2100 for approximately 7 million brownfields premises (under the 93% FTTP model). If we assume a customer connect cost for FTTdp of \$400 per premises (little more than a guesstimate), then we have a saving of almost \$12 billion. This figure is somewhat meaningless as it neglects committed FTTP deployments outside the greenfield FTTP areas, but it gives an indication that the savings under the FTTdp approach are potentially substantial. These savings come with little impact on revenue or operating costs. Should service adoption rates exceed the 70% used as a conservative estimate, the FTTdp business case improves even further. The marginal cost of connecting additional customers to a FTTdp network is very small.

A \$12 billion reduction in costs provided by this back-of-the-envelope analysis certainly improves the outlook of both Scenario 1 (FTTP) and Scenario 2 (Radically Redesigned FTTP) presented in the NBN Co Strategic Review report. However the more sensible approach is to consider the adoption of HFC and FTTB such as starting from Scenario 4 as a baseline. Here the lower number of FTTP premises in Scenario 4 translates to a reduced saving from the FTTdp approach. We can assume that FTTP is still used for greenfields, so the 63% listed for FTTP equates to a total of approximately 7.687 million premises, of which 5.576 million would be brownfields, and hence FTTdp. Allowing for a 70% connection rate, the savings are \$6.6 billion.

This modified Scenario 4 then looks very favourable in terms of costs compared to all the other technology mix scenarios presented in the Strategic Review document. The network capability benefit, including the vital benefit of network flexibility to handle higher rate CIR services, high demand users, and higher traffic demand profile increases, is significantly improved compared to other scenarios involving substantial FTTN deployment.

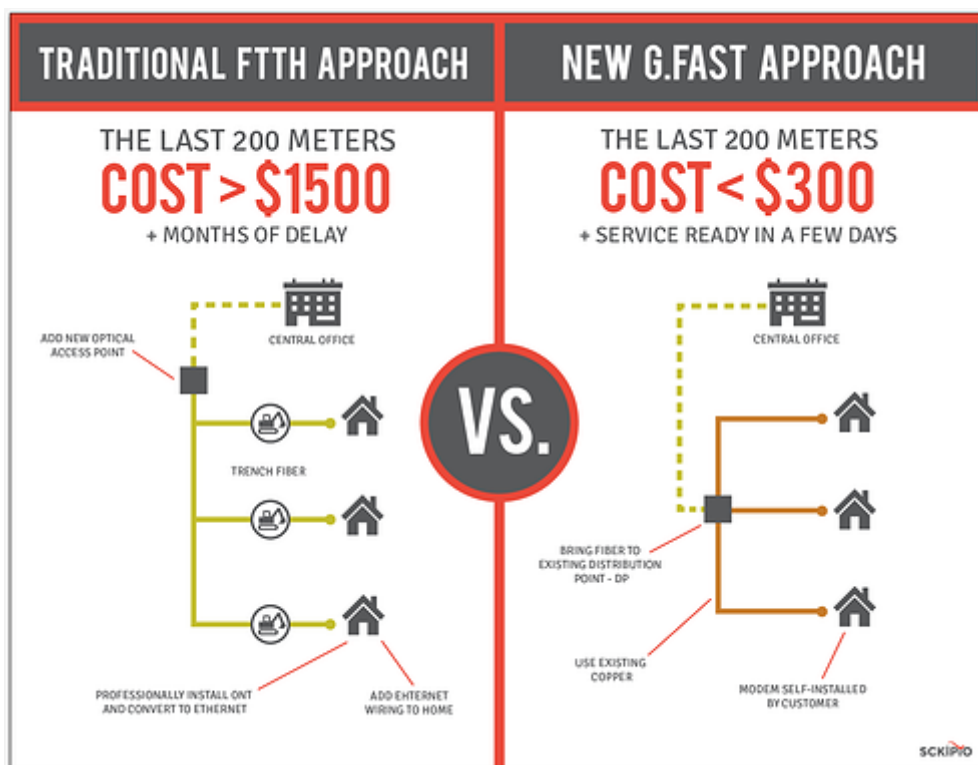


Figure 8: A Cost Comparison from Skipio, [19]

Figure 8 above reproduces a cost comparison graphic from Skipio. The numbers may not be

particularly relevant to the Australian NBN context, but the key point of the graphic is to highlight the fact that the cost differential is substantial. Interestingly, Sckipio is a young semiconductor company focussed heavily (solely?) on the emerging G.fast market^[19].

Negotiating a sensible price on the FTTdp micro node units is a significant step in being able to properly evaluate the real cost savings from the FTTdp approach. Importantly this technology is now seen as relatively mature (for VDSL at least) and multiple vendor options exist.

FTTdp involves one phase of major network build. The capabilities of short-loop VDSL2 and G.fast imply little need for premises fibre drop on a long-term basis except in the more extreme traffic demand growth prediction cases. G.fast upgrades and individual fibre drop connections can be implemented purely on an as-needed (user pays) basis. Another major construction process at a later stage is not required.

G.fast is entirely flexible in upstream versus downstream profile. VDSL is unsuitable for deployment where crosstalk exists between neighbouring pairs using different profiles. This effectively limits FTTN VDSL to the provision of a “one size fits all” selection of profile. FTTdp provides an option of escaping this single global profile. However, the significantly increased average connection bandwidth in the FTTdp setting ameliorates the profile restriction concerns in most cases.

The variable cost of additional premises connections is minimal once sufficient connections have been made to ensure that all micro-node devices are installed. The implication is that any gains in subscriber numbers beyond the 70% figure assumed in the NBN Co figures comes at minimal incremental cost in the FTTdp case. A similar reality applies to FTTN, but the more restricted bandwidth offering of FTTN limits the prospect of heavier subscriber adoption. In the FTTdp case additional subscribers contribute more to profitability than to cost. Given the uncertainty of the 70% figure, this represents an upside potential that ought to be taken into consideration at an appropriate level.



Figure 9: A small pair bundle is enough of a headache!
(extracted from [17], AFR)



Figure 9a: Is a large copper bundle a migraine or a stroke?

Figure 9 above shows the amount of confusion present in the copper network even with a very small feeder bundle. In addition to obtaining extremely high VDSL performance, the FTTdp approach eliminates upwards of 90% more of the copper network compared to FTTN (on a lineal

pair measure basis). The elimination of the vast majority of copper pair connections is a major advantage even if we assume that there is no major network-wide corrosion problem with poorly sealed joints or a major issue with faulty gel sealant material. (Figure 9 was extracted from [17], Australian Financial Review.)

Passing the cost of CPE more directly on to end users, most logically through RSPs, is unlikely to meet with significant objection under the FTTdp model. VDSL modem equipment is expected to be readily available at reasonable prices from a large number of suppliers as is clearly the case today for ADSL modems. (The same claim can be made with FTTN, but CPE equipment costs for FTTP present some concern with how these are likely to be passed onto customers and not act to slow service uptake.)

Cut-over from a PSTN/ADSL connection to an NBN FTTdp connection can easily be accommodated by a single person mobile unit needing to work only at the pit location. Workloads in any one area can be balanced by scheduling, with new CPE posted or delivered to end-user premises ahead of the scheduled time of disconnecting the old CPE and connecting the new equipment.

Micro-node devices, certified by NBN Co for use on the NBN, can be installed and tested as part of the FTTdp fibre deployment or on an as-needed basis for connections. Both approaches have pros and cons, and an optimised approach may involve a mixture. **In either case the skill set required for premises connection is modest and mobilisation of a large workforce of individual contractors, certified by NBN Co to work on the NBN and sponsored and managed by RSPs, is a realistic expectation.**

The premises connection process being largely managed by RSPs has some advantages to NBN Co, although the detail of how the NBN Co/RSP interactions will occur needs to be determined. Of particular note is how individual contractor efficiencies can be maximised by limiting unnecessary geographical movements where the installation process is controlled in a distributed fashion over a number of RSPs.

The FTTdp approach has the potential to reduce LFN fibre counts and substantially reduces the size of splitter cabinets required (FDH enclosures). This opens the prospect of splitter infrastructure being retracted into pits, away from vandals, errant vehicles, and falling branches. The reduction in impact of street furniture is a major positive consideration compared to the FTTN approach. Large cabinets present a very real site location problem that altered legislation does not 'solve'. NBN Co and the government can not expect to win the hearts and minds of the masses by treating genuine public concerns with contempt.

The benefit of reduced LNDN fibre counts is unclear in terms of impact on deployment costs. Presumably smaller and lighter cables are somewhat easier to pull or blow through conduits. Small incremental differences in costs can translate to large savings over the size of the NBN project. The complexity of the task of accurately patching LFN fibres at the FDH is reduced dramatically when there are only 50 or so fibres to contend with compared to 600.

A FDH serving an average of 200 premises requires 600 fibres on the downstream side of the splitter if we assume a per premises fibre count of 3. These fibres are carried in various directions from the FDH in bundles from 72 to 288 (using 12 core ribbon fibres) under the current network design rules^[5]. The NBN Co Strategic Review document suggests an option with a radically different number of fibres per premises, so presumably there is a significant saving to be obtained via a reduced fibre count. On the upstream side of the splitter (the DFN part of the LNDN), there is effectively no difference in fibre count requirements in the FTTdp scenario compared to FTTP.

On the downstream side of the splitter (LFN) there is the potential to use loose tube fibre bundles from 6 to 24 fibres (corresponding to the 72 to 288 fibres in the initial FTTP approach). A single fibre goes to each micro-node device. In reality some additional fibre capacity would be included in the bundles. For FTTP upgrade the micro-node device is replaced by a small splitter (a dual-stage optical splitter FTTP implementation thus results). Micro-node devices with integrated splitters are clearly not difficult to envisage.

The small size of a connector for a single fibre compared to a 12-port multiport enables some flexibility in terms of the LFN deployment. Presumably a major component of the cost in the LFN deployment is not pulling the fibre itself, but ensuring that multiports present at the right locations in the pits. If field assembly of a unit can be performed prior to pulling or blowing through conduit, then the operation is far simpler than the alternative of feeding multiple multiport tethers to centralised splice locations. The relative flexibility of loose tube fibre provides the ability to cater for field uncertainty in the lineal conduit measurement between pits, and provides an ability for safe work practices should tapping into a spare fibre in a bundle become a later requirement at any particular pit location. With sufficient flexibility fibre bundles can be carefully coiled (perhaps using a holding frame) at each pit location without exerting fibre-destroying torsion, allowing sufficient 'slack' at each pit for meeting both objectives above.

Suggestions of major potential LNDN deployment cost savings from FTTdp are speculative, but warrant close examination. NBN Co is able to supply the detailed cost information required to properly delve into such matters and should undertake a comprehensive analysis.

6.3 Additional Cost-Benefit Considerations

While there is solid reason to suggest that a primary deployment of FTTdp compares very favourably against the FTTP or FTTN alternatives, we must also consider some of the not-so-obvious benefits and risks involved with FTTdp.

FTTdp risks supplying greater capacity than required in the short term, but this risk has negligible impact for two primary reasons. Firstly we note that more capacity is required than what is provided on average currently, and hence if we do not choose the FTTdp option we must implement another option such as FTTN. FTTN has limited risk of supplying greater capacity than required in the short term, but given the cost equation is likely to be similar to FTTdp, this is far from being a preferable option. Secondly, given the bandwidth demand arguments presented above, it is unlikely that FTTdp will provide any significant bandwidth oversupply. Rather, the capacity promised by FTTdp is an extremely good match for a broad range of what appear perfectly reasonable bandwidth demand projections.

The risk that FTTdp might under-supply bandwidth is real (on a slightly longer term outlook basis). However, the flexible upgrade nature of FTTdp to FTTP on an ad-hoc basis (and the promise of G.fast) makes the consequences of this risk minimal.

A very tangible but difficult to quantify benefit from primary FTTdp deployment is that it builds telecommunications infrastructure skills bases that will serve the nation's long term future. There is

far less direct connection between the skills needed for FTTN deployment and those required (most likely within a few short years) to upgrade to FTTdp or FTTP. A large part of the army of installers required to roll out the FTTdp network will be able to find ongoing work in the case where there is a growing demand for individual G.fast or FTTP connection upgrades. This provides long-term stability to the work force, in turn encouraging work force growth in the crucial early network implementation years.

Crucially, the requirements of high-demand users are met by the FTTdp build to the same level they are met with FTTP. We must not be tempted to think that high-demand users represent only a small percentage of the total user base of the NBN and their requirements should only be weighted in proportion to their number. The NBN represents important national infrastructure due to the fact that communications networks are inextricably linked to the modern global, and rapidly evolving, economy. A significant proportion of high-demand users are likely to be involved with ventures that represent a growing component of the national economy. High technology web companies are part of this equation, but greater use of networks is implicit in almost every sector of our modern economy.

Business investment in smaller organisations can be expected to be predicated on reliable network connectivity that promises to grow with the business. A fibre-on-demand model for individual FTTP connection in a FTTN deployment meets some of the generic requirements of business customers, but it presents some uncertainties related to cost and installation time. While these issues do not pose a major threat to many businesses, they are likely to be significant concerns for others. The FTTdp model provides baseline capability that will suffice for the majority of customers, and simple upgrade options. As such it instils communications capability confidence for businesses, providing a communications technology roadmap that can be relied upon.

The prospect for uniform retail product offerings across the breadth of the NBN is significantly greater with FTTdp than with FTTN. The extra complexity of not being able to supply a largely uniform retail product set has business cost implications for RSPs.

For many premises the high utility afforded by FTTdp or FTTP is only likely to be fully exploited by video streaming, largely for entertainment purposes. In the FTTP case it appears that the cost differential compared to FTTN is substantial. A detailed cost-benefit analysis would thus likely be required to select the more sensible and strategic approach for the nation if only comparing FTTN and FTTP. The FTTdp option promises similar overall capability to that of FTTP at a significantly reduced cost, perhaps even at a reduced cost compared to FTTN. A cost-benefit analysis is hence likely to weigh heavily in favour of FTTdp.

7. Large and Small MDUs

MDUs account for a significant percentage of the nation's premises. Larger MDUs (those with greater than 30-40 premises) are noted to account for approximately 1 million premises in the NBN Co Strategic Review document. The difficulty in routing new fibre optic cabling internally within many MDUs is a reality that has plagued NBN Co to date. (Difficulty in routing new coaxial cable in MDUs caused related difficulties for HFC network construction efforts in decades past.) Part of the difficulty is that practical solutions such as options for conduit placement may not meet with body corporate consensus. A completely different approach where MDUs are responsible for their own connections is suggested.

Paralleling the FTTdp model, NBN Co should retain responsibility for bringing fibre to the building/basement. The default NBN Co option should be connection of powered node equipment at this point, implementing a FTTB solution. If prior to the initial NBN Co connection the body corporate has engaged private fibre cabling contractors (NBN Co approved suppliers) and produced fibre cabling to the basement or other appropriate communications enclosure, then NBN Co should supply a FTTP solution to MDUs.

Crucial to this approach is that NBN Co approved private contractors are responsible for all in-premises installation and all on-site design and installation, leaving NBN Co to provide the fibre interface (splitter or patch panel) only. While there should be no compulsion for individual premises to install CPE (a futile exercise if there is no prospect of an NBN Co service being ordered), NBN Co can insist that fibre is provided to the entire MDU to the extent that cabling is installed in risers ready for easy deployment at a later date. There is little marginal cost benefit from provision of fibre capacity to only 70% of MDU premises that wish to subscribe to NBN Co services, but unless enforced, private contractors will take short-cuts.

Such an approach fosters a healthy ecosystem of private contract installers attempting to find winning designs to satisfy the demands of body corporates. A default connection over in-building twisted pairs is provided, and the body corporate is responsible for anything more, including dealing with the limitations of decrepit copper infrastructure should this exist.

Vectored VDSL in a FTTB setting can be expected to deliver robust connections. However, a small percentage of premises are likely to demand more than what FTTB provides. Over time this small percentage will undoubtedly grow. The economics of provision of FTTP on an individual basis in an MDU setting might work out cheaper than cabling the entire building (certainly the case if the particular premises is close to the communications cabinet), but the added complexity from the NBN Co side of dealing with a mixed FTTB/FTTP solution also needs consideration. Encouraging MDUs to investigate provision of their own fibre cabling prior to initial NBN Co network installation can ameliorate this effect, as those MDUs more likely to require individual FTTP connections at a later stage correlate strongly with those likely to install their own building-wide cabling.

NBN Co should give consideration to the option of MDUs being provided with control of their own circumstances (FTTB versus user-pays FTTP), fostering a private (independent) contractor ecosystem.

Vectored VDSL solutions do provide a limitation on upload versus download profile. In the short copper loop setting of MDUs this is unlikely to be a major concern for the vast majority of users. However, MDUs must be as fully informed as possible and supported in their decision making process with regard to the prospect of private fibre cabling prior to NBN Co connection. Comprehensive web-based content can be provided to MDUs both generally via the NBN Co website, and with a link provided at time of early NBN Co direct contact to body corporates (there is a need to satisfy informed consent provisions both from a regulatory perspective and a moral perspective). The prospect of G.fast deployment in the MDU environment is yet to be fully understood (as touched on above). This fact must be highlighted and information provided as it becomes available.

For some larger MDUs it is conceivable that existing in-building cabling other than twisted pairs can be employed for the access network connection. The simplified options are twisted pair versus fibre, but where coax infrastructure exists there is a need for more detailed consideration. NBN Co may need to develop clear guidelines if this is likely to be an issue for a non-negligible proportion of MDUs.

There is undoubtedly some complication added on the NBN Co side related to having heavily mixed regions of FTTN/FTTB and FTTP. A full FTTP deployment naturally has no such complication. Widespread FTTdp deployment likewise promises minimal additional complication due to the flexibility of the implementation to incorporate on-demand FTTP. For a mainly FTTN deployment, the desire to upgrade a number of connections to FTTP must be accepted. It is a complication that can not simply be relegated to the too-hard-basket. Equipment vendors appear to be recognising this by supply of node devices incorporating optical splitter capability and designing enclosure solutions based on the assumption of the need for future splitter installation.

7.1 Small MDUs

Small MDUs also account for a sizable proportion of the nation's premises. It is arguable whether it is any easier to obtain agreement from the body corporate of a small MDU compared to a large MDU. Small MDUs perhaps also account for a greater variety in construction, and the difficulties of finding suitable fibre designs is magnified on a per premises basis compared with the potential scalability of a fibre design in the large MDU setting.

One of the authors has direct knowledge of a small MDU in Brisbane of 8 units that has recently been through the NBN Co design process. The complex has a contiguous roof line linked by carport areas beside individual units. With the exception of one second floor unit, all other units are ground floor construction. Moderate slope to the ground level at the site compounds the challenges of installing fibre cabling through the roof spaces. Ultimately multiple site visits and a significant amount of time from both designer and a rod and rope crew have been required, just to produce a design proposal for the body corporate to consider. It doesn't take much imagination to see that this level of complexity is not encouraging for widespread FTTP deployment to smaller MDUs.

Under the FTTdp model, such a MDU can be served directly from a pit-mounted micro-node in the street or in the MDU's Telstra MDF enclosure. The latter option is obviously preferable, and comes with little cost implication if a fibre tether can be easily installed using existing conduit. For the above-mentioned example, such use of the existing lead-in conduit is already part of the NBN Co FTTP design. Installation of a micro-node device at the MDF location is thus a comparatively simple process compared to FTTP deployment in the small MDU category. We can go so far as to state that there is little reason to suggest that FTTdp deployment for small MDUs would involve additional difficulty compared to SDUs.

FTTB with the option of MDU-funded FTTP appears to make sense as the paradigm for larger MDUs. Smaller MDUs should be treated in much the same way as SDUs in the FTTdp model. In areas where FTTP is otherwise used, micro-node deployment for small MDUs must be considered to parallel FTTB for the larger MDUs.

7.2 Mobilisation of a Contractor Ecosystem for MDUs

One model of NBN deployment is central control of work palmed out in detailed contracts through a hierarchy of contractor companies. There are many reasons to prefer such a model, and it serves the nation well over the broad deployment of the NBN, whether that is construction of the transit network, a FTTP access network, a FTTN network, or FTTdp.

However, there are also reasons to desire a rich ecosystem of individual contractors competing to provide winning designs and installations to non-centralised decision makers. While there is limited scope for this in the individual premises network connections, except via the avenue of multiple RSPs, MDUs do show significant promise for substantial contractor autonomy. We have already identified above the potential for MDUs to engage individual contractors to provide competitive in-building fibre designs and installations. We ought to consider if a slightly broader ecosystem of providers should be encouraged.

While today there would seem little to be gained by suggesting micro-node devices could be installed in risers between floors in larger MDUs, this 'radical' idea might be of interest in a minority of cases at some future point. NBN Co should ideally provide a specified group of fibre service options to the building/basement, with equipment, designs, and implementation from there being controlled by the body corporate or building owners in conjunction with a healthy supplier ecosystem. NBN Co can provide clear specifications and guidelines, including approved equipment options and registered suppliers.

Should a particular MDU's needs or desires evolve, then they ought to be able to upgrade in-building infrastructure as required. In some cases this will clearly also involve modifications on the NBN Co side, and a schedule of charges may be applicable. With an ecosystem of suppliers established, the best possible outcomes can be obtained with regard to re-use of decommissioned equipment. We can consider parallels to the decades of history in the used PBX/PABX market.

Building owner and body corporate autonomy, supported by a healthy ecosystem of designers, suppliers, and installers, is great for the NBN and good for highly skilled small business operators. Such a pool of individual operators are likely to overlap with the FTTdp installer pool, ultimately creating a large army of adaptable technicians able to flexibly supply resources as required. By encouraging opportunities for smaller operators we may even reduce the prospect of skills loss after major construction phases are completed.

7.3 MDUs in the HFC Envelope

Large unserviced MDUs in the HFC coverage areas require careful consideration. Deployment of new cabling of any sort (coax, Cat 6, fibre) within these buildings involves a significant effort. We must thus firstly ask whether the FTTB model, exploiting the telephony wiring, is the better option. This requires fibre to be deployed to the site, but under the assumption that large MDUs are not isolated buildings, the fibre distribution network to the building is a shared investment. We thus expect the economics of FTTB to be sound for a significant number of MDUs in HFC areas.

Should FTTB not prove economical due to the cost of fibre deployment to the building, then new fibre cabling within the building is a wiser long term investment than coax deployment. The provision of RFoG (RF over Glass) becomes a sensible approach, and for larger MDUs it is not too difficult to imagine the HFC fibre being brought into the building's communications room. It is understood that these considerations are not unique to the Australian NBN setting, and that plug-and-play equipment alternatives exist in the cable networks' world.

Another option that must be considered is that of installing Cat 6 cabling within the building. We can imagine that Cat 6 cabling might be more tolerant of sharp bends and tricky physical installations than fibre. With ethernet supplied over Cat 6 from the basement or communications room, the CPE side of the equation is clearly simplified. However, the desire to present multiple UNI-V and UNI-D ports (User Network Interface, Voice and Data) to the customer in a uniform fashion over the NBN would seem to require more analysis.

The complication that presents here is the chance that MDUs are treated differently depending on whether they are in HFC network areas or not. From above it was suggested that outside HFC areas NBN Co is responsible for delivering a small selection of fibre connectivity options to the building basement. The building owner or body corporate is then responsible for their own installation. The fall-back or default option is use of the existing in-building copper pairs for FTTB. In the HFC areas if we eliminate this FTTB default option then we do not supply building owners with a minimal cost solution. The level of negative sentiment generated by such an approach might be expected to be significant even with a very small number of affected MDUs.

If we adopt a model where MDUs are entirely responsible for their own in-building installation, including the cost of any node equipment, if needed, then the marginal cost for the affected MDUs in the HFC footprint areas is reduced. However this comes at the cost of a large number of MDUs having to levy individual unit holders to fund even the default connection option. Clearly this will also produce a significant amount of negative public sentiment.

The preferable option is perhaps for NBN Co to fund in-building fibre cabling for those MDUs in HFC areas where NBN Co opts not to supply fibre to the building. Detailed studies may reveal that only a very small number of MDUs are likely to be impacted by this issue in any event.

The cost of funding such in-building cabling and the cost of adding additional CMTS/CCAP capacity and any node splitting involving deployment of additional underground or aerial coax, must be weighed as part of the NBN Co decision on whether to supply a fibre connection (FTTB) to the building.

8. Mobile Connectivity

The NBN is perhaps rightly fixed-line centric. Our modern communications network experience incorporates both mobile and fixed-line connectivity. While some have suggested that our entire communications demands can be supplied by wireless, the physical reality of a limited spectrum implies that wireless and fixed-line technologies will co-exist indefinitely. They are correctly seen as complementary technologies in many ways.

In recognition of the supreme importance of mobile device connectivity to modern life, we must ask whether there is any direct way that NBN Co physical infrastructure deployment or service offerings can be tailored to better serve the needs of mobile connectivity. Bill Morrow, as CEO of Vodafone Australia (and recently appointed to the role of CEO with NBN Co) suggested NBN Co ought to consider this issue in his statement introducing the McKell Institute report “Superfast Broadband: The future is in your hands”^[12] released in October 2013. Indeed, the prime recommendation of this report is that the NBN should seek to consider how it can improve mobile connectivity outcomes.

It must be noted that while there is no doubt that the long-term interests of the nation are well-served when available public funded infrastructure is shared on equitable terms to all enterprises, there are vested commercial interests that make the interplay between the NBN and wireless operators a complicated matter. Without an NBN Co infrastructure deployment, mobile carriers face a costly process of investing in back-haul capacity. In some cases installed dark fibre exists, but back-haul pricing is influenced more by the cost of duplicating such capacity than by the marginal cost of using spare capacity. In the commercial environment of a small number of network owners such is only natural.

The risk posed to mobile network operators is a possible disruption to their future income from rental of back-haul network capacity to third-party operators, or through loss of customers to such operators that are able to compete on a leaner financial basis having avoided significant investment in back-haul capability. **Nonetheless, NBN Co should consider what long-term benefits might be afforded to the mobile communications industry.** Pleasingly, the early signs are present that this is precisely the path NBN Co is taking. A process is under way to trial provision of back-haul service to mobile phone towers^[13]. The technical considerations here are whether the NBN network can supply services that are adequate for the task. The mobile back-haul requirements differ sufficiently from the primary NBN design requirements to suggest that a trial will obtain useful information and increase the NBN Co knowledge base in this area.

8.1 Exploitation of NBN Assets for Mobile Augmentation

It is recognised that site locations of NBN Co fixed wireless towers are not obviously suited to cellular mobile approaches, and that the use of microwave point-to-point links of only moderate capacity provides limited back-haul connectivity to towers in many cases. Where fibre deployment was planned as part of FTTP or is likely to occur to support FTTdp, or FTTN, the chances are that multiple cellular networks already operate.

With rapid growth in data demand over cellular networks, and the possibility of large demand spikes occurring through OTT (Over-the-top) services, greater use of HetNets, micro-cells, and

deployment of strategically placed Wi-Fi will be of interest to mobile carriers. The NBN is likely to have a significant role to play in this area, but to some extent this is already catered for through the inclusion of fibre capacity for non-premises deployment.

The organic Wi-Fi network approach where a common SSID is present on all Wi-Fi routers connected to a service provider's access network is perhaps a bit too radical to suggest it will receive widespread deployment. However, there are already Wi-Fi mobile telephony 'carriers' such as US-based Republic Wireless offering cheap mobile call plans with no cellular network service.

NBN Co must be flexibly placed to meet any service demands introduced by emergent technologies. We must not be blinded by paradigms of the past, and ought to investigate the mobile connectivity issue very carefully.

8.2 Caravans and Motor-homes

At any one time a small percentage of the nation's population resides in caravans and motorhomes. Caravans obviously cover a broad spectrum of use. We would question what level of benefit the domestic population who choose an occasional caravan trip would get from improved broadband access solutions, but other usage scenarios stand to benefit significantly. Reliable broadband access to caravans and motor-homes would potentially stimulate large areas of the tourism market, domestically and for overseas visitors, where there is a combined thirst for travel and a need for continuous connectivity. Some people may be enabled to provide on-going business activity as they do the 'grand tour' of the nation.

NBN Co should consider whether the planned satellite service offering can be tailored to the needs of such a niche market. Alternatively, the Department of Communications may wish to investigate how these customers might best be served. It is a lofty goal to provide superior broadband connectivity outcomes everywhere, but the potential benefits are significant once we consider global trends of increasing daily reliance on connectivity. Perhaps Australia could be seen as a tourist destination where you can get away from it all while never losing contact!

9. The NBN Deployment Prioritisation Issue

There are several competing considerations that enter into the equation when determining where network construction resources should first be deployed. A primary objective must be to maximise the investment outcome. However, spending X dollars to take 1000 customers from good ADSL2+ service to the NBN must be weighted carefully against the same X dollars to provide NBN service to a smaller number of customers limited to sub 256 kbps connectivity.

Pragmatically, there are many elements that need careful weighing. Workforce location factors and the need for the general public to be able to understand the rationale behind prioritisation decisions are just some of the considerations. Given the long construction timetable, it is important that deployment prioritisation decisions are transparent and not open to abuse of any sort.

Creation of a multi-interest coordination committee to oversee the NBN deployment decisions is perhaps necessary. The federal government's broadband survey would provide valuable input to this committee. Such a group might have representation from the Department of Communications, NBN Co, and major existing industry broadband suppliers. The terms of reference must be precise and act to limit the very real potential for conflicts of interest, keeping the best interests of the nation at the forefront. An extremely high level of transparency, with fully public minutes of all meetings can help ensure this objective is met.

Crucially, current broadband network service providers can make sensible investment decisions on where to supply upgraded network equipment over the long NBN build phase by having reliable knowledge of NBN deployment schedules. The paralysing effect that comes from a lack of such knowledge creates negative communications network outcomes that our nation simply can not afford. If the NBN is known to be 6+ years from a particular location then commercial decisions about equipment and service upgrades in those locations may enable incrementally higher levels of interim service to be provided by the current broadband suppliers.

Whether deployment prioritisation is executed in a committee structure or via an alternative governance structure (perhaps internal to NBN Co and the Department of Communications), it is recommended that all deliberations are fully published as a matter of public record.

There is no strong argument against a very high level of transparency. The public has a right to know why one area might receive NBN service 7 or 8 years ahead of another. Many communities feel the pain of inadequate communications network services and this will continue to be the case until the NBN build is complete.

The committee structure approach has the advantage of eliminating the perception of political meddling in the deployment prioritisation task. This creates an environment of long term stability (looking beyond the horizon of the next election) that is crucial to invigorate commercial investment. There is also a benefit to the all-important public perception of the NBN. While this last point is difficult to quantify, it is not insignificant.

10. Conclusion

- To the extent that FTTP proves to be prohibitively expensive and time-consuming to install, DSL technologies must be adopted.
- FTTN receives justifiable renewed consideration with the commercial deployment of Vectored VDSL solutions. VDSL profile 17a with full vectoring appears to promise up to 50 or 60 Mbps download at a copper distance of up to 800m from the node (although it remains to be clarified whether this performance can be obtained in the Australian twisted pair copper network context). The uplink for the same 17a profile at this distance is a more modest 15 Mbps with full vectoring implementation. (Of crucial note is that bonding and phantom mode are only applicable where more than a single pair is available.)
- Such data rates are promising when assessed against the user demand known to exist today. Only a small fraction of highly demanding users presently exceed these requirements. Allowing for potential demand growth over as little as the next three years the equation becomes substantially less clear.
- A deeper fibre penetration into the access network must be fully investigated and properly weighed against the alternatives. FTTdp promises to meet the vast majority of all demands over the next decade or more with short-loop VDSL2 implementation, and has straight-forward upgrade paths to G.fast and FTTP on an as-required basis.
- If we start from a baseline assuming the use of HFC, similar to Scenario 4 in the Strategic Review, the saving from FTTdp (based on back-of-the-envelope analysis) is approximately \$6.6 billion. This figure is substantial and places the cost of widespread FTTdp deployment in the same ballpark as that of FTTN. More detailed analysis is obviously needed.
- While the cost equation is appealing at a back-of-the-envelope level, a decision to pursue a FTTdp trajectory versus a FTTN one must clearly be made on the basis of a consideration of both costs and benefits. Here the network utility provided by FTTdp is a large improvement over that provided by FTTN. The FTTdp network capability level is a clear winner in most user demand projection cases.
- FTTdp has the advantage of eliminating somewhere up to 90% of the legacy copper network compared to FTTN. Importantly the majority of the joints are eliminated.
- Primary use of FTTdp is likely to lead to some revenue upside compared to the FTTN case, due to the higher connection rates and the greater ability to offer high CIR products.
- The revenue upside could be substantial in FTTdp should subscriber numbers be driven upward. While this is true in FTTP, it comes at additional large customer connection costs that simply do not apply in the FTTdp case.
- FTTdp and FTTP are closely aligned, and there are a number of expected operating expenditure savings to come from having FTTB but limited FTTN deployment in the network.
- A more sophisticated analysis such as that performed for the other technology options in the NBN Co Strategic Review report may indeed determine that widespread FTTdp is the

obvious technology choice from all angles, chiefly cost, time, and capability (benefit).

- A multi-interest coordination committee (or other transparent governance structure) to oversee NBN deployment prioritisation could ensure overall broadband services provided by all current broadband providers are maximally exploited to meet customer needs during the lengthy NBN build stage as well as assist NBN Co to roll out the NBN in stages to meet most urgent needs first.
- HFC networks clearly have capabilities to provide high levels of NBN connectivity. With an appropriate technology road-map it is possible to envisage HFC networks providing high service levels for at least a decade and perhaps substantially longer. This potential must be fully exploited by NBN Co.
- Detailed engineering analysis is necessary to produce a viable technology road-map for HFC network areas. Compared to other access technologies, the multiple legacy capabilities of HFC networks and the rich global industry environment, makes the provision of a thorough technology road-map a complex communications systems engineering undertaking (with due attention to business case aspects). It is a process that needs to be commenced without delay.
- NBN Co certified mini/micro/basement DSLAM equipment, lead-in fibre cabling and building fibre cabling could be employed at the customer end of the NBN and these could be installed by NBN Co certified independent contractors/technicians sponsored by RSPs, etc. This appears to be particularly relevant in MDU settings to empower building owners and body corporates, in the process establishing a healthy ecosystem of independent contractors.
- Fixed wireless should only be deployed where VDSL implementations are unlikely to be viable. The cost-benefit comparison between VDSL and fixed wireless comes down strongly in favour of VDSL in settings where clusters of premises exist. This is especially true when we allow for the reality of evolving premises bandwidth demand statistics.
- A detailed and honest analysis of bandwidth demand projections, including identification of temporal and individual user variability, encompassing a range of bandwidth growth projection scenarios, will provide valuable input to network provisioning considerations at many levels within the NBN.

Polarised views will continue to plague the NBN initiative while FTTN is a primary deployment technology. We can expect that within as little as 3 years one side of this debate will be able to point to enough external developments to quieten objection from the other side. The risk in FTTN is that we may be on the wrong side of the equation. FTTdp holds the promise of eliminating the vast majority of any polarised debate from square one (and may do so with no additional cost implications). The nation must honestly investigate this option.

Appendix A: Contention Basics

Metrics such as “Contention Ratio” and “Average Busy Hour Throughput”, (ABHT), are easily calculable and often used when performing network engineering design tasks. They are employed as a proxy for understanding network utility, providing guidance on the need to allocate additional link capacity where desired. Choosing between closely related design options with an increased ABHT or decreased Contention Ratio, translates to an undeniable increase in network utility. However the relationship between such metrics and the all-important factor of network utility is extremely non-linear in a very complex fashion. As engineering decision makers it is easy to become confused about the complexity of this relationship to utility, especially when we are accustomed to regular use of the metrics in contexts where the relationship to utility is more mundane or can be assumed to be a simple linear relationship.

We must firstly appreciate that the relationships between the metrics and network utility is highly non-linear over a large range of values. A doubling of ABHT from 0.5 Mbps to 1.0 Mbps does not necessarily produce the same ratio of difference in network utility that doubling ABHT from 10 Mbps to 20 Mbps produces. Our common interest in such network metrics does not usually cover large ranges, but the NBN context of selection of drastically different network access technologies, very large differences can appear. For instance under the current NBN Co design philosophy, a FTTP area is provided with an initial ABHT of 10 Mbps, while the long term satellite solution provides for ABHT of 350 kbps. Extreme care must be applied in this context, as will be detailed below.

Less obvious is the fact that the relationship between metrics such as Contention Ratio and ABHT and network utility evolves over time. This is a simple consequence of the evolution of user demands over time. However, the ramifications are far from simple, as user demand is not characterised by a single-dimensional parameter which increases over time. ABHT of 1 Mbps today may represent less utility than ABHT of 256 kbps 10 years ago. **We must avoid the self-congratulatory impulse to view an improving metric over an extended period as positive while ignoring the evolution in the relationship between the metrics and network utility.**

While we see monotonic increases in network traffic, user demand driving this is evolving in disparate directions for different individuals. Some users undoubtedly remain primarily engaged in web browsing, email, and Facebook. Richer content on the average website and higher resolution photos uploaded to Facebook, coupled with increasing usage times, drive demand growth for such users. However, a growing class of users are earlier adopters of more bandwidth-intensive applications. It is not crucial to delve into the specifics of what these applications may be at this point. The crucial observation is that while demand profiles have always varied between users, the extent of this variation has increased dramatically in recent times compared to 10 or 15 years ago.

The use of any statistical averaging of usage must consequently be carefully considered. It is one thing to discuss ABHT averages on backbone network links from a large exchange dealing with perhaps as many as 10,000 subscribers in a major European city, but it is quite another to translate this metric to access network sharing of resources between a small number of users.

A.1 Shared Links



Figure 10: We all need to share! (Image Credit: Datacom Systems)

Shared links are an obvious requirement of any non-trivial network. When such links are in the fibre-optic backbone network, there is a solid economic case for reliable provisioning of bandwidth to meet not just the average demand, but much of the peak throughput demand as well. In the context of fibre-optic backbone provision for the voice-based PSTN of the past, the ability to meet all peak demand was feasible (although perhaps not practical). However, each user in the PSTN example can be assumed to have a modest 64 kbps steady demand. Today peak demands are up to three orders of magnitude greater (depending on access network technology limitations), and individual users have extremely bursty data use profiles. While it is likely to be prohibitively expensive to provide enough shared bandwidth to meet the worst case peak demand, it is possible to economically provide capacity to meet all but a very small number of cases. Provisioning may adopt a 3-sigma, 6-sigma, or any other design metric target, depending on application demand. Where traffic is growing steadily, metrics can be applied to track the changes and provision capacity in timely fashion to ensure overall network transport criteria are met.

Adaptive provisioning of capacity becomes a more tricky proposition where network traffic growth is less predictable. The most obvious factors at play here are the increased availability of 'new technologies' such as high-power mobile devices, cloud storage, and emerging 'apps'. Any one or any combination of factors can lead to sharp demand spikes that may be foreseen too late for orderly capacity provisioning lead times.

Dealing with such factors is the life-blood of any network or communications company, and is closely linked with company costs, revenues, and establishment and protection of the brand. It is perhaps easy to point to examples where various companies have excelled in this crucial area. A few companies have made major blunders, causing failures, and most entities manage this area 'adequately'.

Contention becomes a far more vexed issue with something like the NBN due to the variation in factors between different access network technologies. The decision about which technology to employ in any given setting can be crucially impacted by a proper analysis of contention.

A.2 Statistical Usage Patterns

In order to perform network provisioning calculations we must make assumptions about individual usage patterns. Network traffic demand varies dramatically from user to user and from one period to the next for single users as application profile changes. Examples are web browsing versus streaming video content. Adoption of simple usage profiles requires the "Law of Large Numbers" for validity. We are able to employ standard statistical methodologies such as Poisson Distribution analyses, to significant benefit where such assumptions can be deemed valid.

In the NBN FTTP context, a Fibre Serving Area (FSA) has a nominal premises count of 4000. In such a setting it is generally safe to assume the use of statistical methods. The use of a 40 Gbps effective link from the FSA to the transit network results in a total capacity limitation (40 Gbps down and 40 Gbps up). This translates to a guaranteed average capacity of 10 Mbps per user, and there is a very small statistical probability, based on current usage demands, of this limit being applied. As usage demands increase, traffic volumes can be monitored and increased link capacity can be added to FSAs on an as-needed basis.

Within the FSA there are other shared network links in terms of individual GPON fibres. In the NBN context a 32-way splitter shares a 2.488 Gbps downlink and a 1.244 Gbps uplink between up to 32 users. This translates to a guaranteed minimum per user of approximately 78 Mbps down and 39 Mbps up (somewhat more if the splitter output is only partially populated). At present usage rates it would require very heavy usage from all 32 users on this single GPON fibre for this limitation to become significant, and the reality is that the above 10 Mbps per user constraint is likely to be significant well before the GPON limit provides noticeable impact. We can understand this intuitively by realising that 85 Mbps down is approximately the bit rate required for streaming 8K video into each premises, or streaming 3 or 4 channels of 4K video to each premises. **While this sort of demand is entirely reasonable to project for a small proportion of premises, the outlook for the next 5 years at least, is that such demand will not be significant enough to cause GPON saturation.**

The conclusion then is that the GPON FTTP approach is well-provisioned to meet current and any future demands. Should usage statistics continue to grow dramatically, at some point the 10 Mbps limitation for the transit link connection may become significant. For the immediate future it represents a guaranteed minimum, but in reality this limit is unlikely to be tested on all but the most extremely rare occasion. An NBN service over fibre promising 100 Mbps down and 40 Mbps up, is consequently likely to deliver this. Limitations beyond the control of NBN Co (in the broader internet) are expected to be the only significant limitations experienced by the end users.

(GPON has well known upgrade paths through XG-PON, with plenty of potential for later upgrade as needed. Individual fibres are able to support massive data rates with only the need to upgrade terminal equipment at either end. XG-PON is expected to be economically viable significantly in advance of widespread network demand at this level.)

A.3 The Statistics of 1 Gbps FTTP

The GPON FTTP network is also fully capable of supporting 1 Gbps down and 400 Mbps up, for a low density of actual demand at this level. The statistical nature of demand and the limited call for 1 Gbps delivery for most applications implies that even multiple 1 Gbps customers on a single GPON fibre are unlikely to provide noticeable impact to other users on the same fibre. More constant 1 Gbps delivery is of course a different matter entirely, however the foreseeable penetration for demand at this level is very slight on the 5 to 10 year time frame.

It is entirely feasible to suggest that having various access speed tiers is counter-productive. While we might expect data usage to increase with widespread availability of 1 Gbps service, the inherent traffic demand is unlikely to change materially from a 100 Mbps service offering. Throttling the delivery rate can be considered to be a highly artificial limit. A user-pays scenario suggests that retail plans should continue to be based on data consumption, as should wholesale charges. There is the undeniable complication with regard to peak versus off-peak data consumption, but ignoring this, a portion of network cost relates to being able to accommodate growth in throughput. It is

reasonable to argue that such marginal costs should be apportioned on an equitable (user pays) basis.

Given that the maximum speed makes little difference to the actual statistics of user demand, a 1 Gbps down 400 Mbps up service offering to all users in a FSA can be claimed to deliver such speeds in reality (on a high percentage statistical basis until such time as individual usage statistics change dramatically).

Should the NBN be providing 100% FTTP, then universal provision of 1 Gbps service is a very real option. There is however a consideration that with a mix of technology options, providing the psychological advantage of declaring universal 1 Gbps service on FTTP would accentuate the impression of a strong digital divide. The extent to which this might be true depends entirely on the real capabilities of the other network solutions employed outside the FTTP areas of the NBN. A FTTdp solution where fibre is deployed down the street and connected to the premises by small node equipment in the pits attached to the twisted pair copper lead-ins, is entirely capable of matching the majority of the FTTP promise. FTTN falls short of this ideal, but shows promise with vectored VDSL2 and shorter copper loop lengths.

A.4 Contention Issues with a FTTdp Network Solution

Fibre-to-the-Distribution Point (FTTdp) can be interpreted in a number of ways as there is no single consensus of what defines a distribution point. For our purposes, we consider a model where fibre is taken down local streets (to the lead-in pits generally located outside property fence lines), but falls short of the FTTP objective of fibre connection into premises. The reasons for emphasising this model is that it is likely to provide a long-term solution in the way that FTTN may not. **With the high performance of VDSL2 over short copper loops, and the emergence of the G.fast standard, the need for a fibre premises drop is far from clear.** The small FTTdp node devices do not fully meet the objectives of an entirely passive network infrastructure from central office to the premises, but they are modular, rugged, and powered from the premises. FTTdp thus checks all the boxes that FTTN leaves wanting.

Provisioning of capacity for a FTTdp solution is easily able to mirror that of the FTTP case, and it is straight-forward to see that the same flexibilities exist for capacity growth on contended links. Note the principle limitation of 40 Gbps for 4000 users creating a guaranteed minimum of 10 Mbps is an effectively identical design criteria in both cases. Similarly, the single fibre splitter fan-out equation is largely identical even though in the GPON case we are dealing with a single layer of optical split (in the current NBN Co FTTP design), and in the FTTdp case there is a lower order optical fan-out combined with multiple premises served via each small micro-node device.

In both the GPON FTTP situation, and for FTTdp, long-term bandwidth demand growth can be met with use of additional fibres or through modulation of additional carrier wavelengths. Such scenarios were outlined for FTTP in the 2012 Analysys Mason report^[3].

Importantly, FTTdp can be readily upgraded to an individual FTTP connection on an as-required basis. The short copper lengths imply that the bulk of premises demand, both now and in the foreseeable future, is able to be met with VDSL2 or G.fast.

A.5 FTTN Contention Considerations

Recently we observed comments made by Dr John Cioffi, CEO of ASSIA, relating to the potential

benefits of FTTN compared to FTTP in relation to the lack of contention comparable to the single fibre GPON network^[4]. While there was expert comment on this matter in recent months covered by CommsDay, it is important to consider the issue objectively.

We noted above a guaranteed minimum of 78 Mbps down and 39 Mbps up for our GPON configuration with a 32-way split. Even allowing all users to be subscribed to 1 Gbps down/400 Mbps up plans, realistic usage assumptions made today would translate to a very small percentage of time that a user would be limited to such a minimum rate. There is clearly a non-zero probability that an individual user would be restricted to a data rate less than an arbitrary value such as 300 Mbps, but this probability drops quickly due to statistical effects as this value is reduced. 78 Mbps down and 39 Mbps up is close to the plan maximum speeds for the BT FTTN deployment. The conclusion is that for present usage assumptions the maximum speed of the BT FTTN deployment is identical to a minimum GPON speed limitation, expected to occur in a tiny fractional percentage of instances.

Assuming demand growth to the point of saturation (a perhaps ludicrous assumption as we must allow for progressive network construction activities as demand grows), the GPON system will deliver a constant 78 Mbps/39 Mbps service^F. This assumes the link from the FSA to the transit network has been resized to accommodate this capacity, and that there has been no change to the splitter ratio or active GPON electronics to allow a higher rate system. As this equates to the maximum BT FTTN plan throughput, it appears that there is no question that the GPON system provides higher capacity and greater overall network utility.

However, while the BT FTTN deployment represents a useful reference point, it is possible to envisage a higher order VDSL2 profile being used that provides for greater maximum speeds. It is also easy to tweak a GPON system to provide a higher guaranteed speed. Reduction of the splitter fan out to 16 doubles the minimum speeds to approximately 156 Mbps/78 Mbps. There are many other options on the GPON side to improve performance, and some of these can be achieved on an incremental basis on individual fibre links according to greatest demand. For example a small number of fibres can be re-patched at the FDH (Fibre Distribution Hub) should usage statistics on any particular GPON segment grow dramatically.

The other point to note is that when we speak of contention we generally assume we are dealing with links that have fixed capabilities, with multiple links coming into a node exceeding the capability of the links exiting the node. In the FTTN context we must allow that contention and crosstalk are closely aligned concepts. To say there are no shared links to individual premises may be strictly true, but the reality is that bandwidth is being 'shared' via the noise introduced by crosstalk. Vectoring provides an ability to extract the maximum output from the system, but it is not supreme magic technology that eliminates all the negatives of crosstalk (just routine magic eliminating most of the crosstalk impact).

In conclusion, while a single shared network link by way of a GPON fibre translates to multiple copper links in FTTN, there is little reason to suggest an improvement in contention on the FTTN side leads to an overall improvement in network utility. We can state with high reliability that FTTN is an inferior option to FTTP in this regard. However, this does not conclude that FTTN deployment should not occur. Such a decision hinges more upon cost and deployment factors in conjunction with capabilities provided related to projected usage demands.

^F We can assume that universal demand growth of this type will allow CVC charges to be reduced to a level where the overall CVC impost is relatively unchanged.

References

1. NBN Co Strategic Review, Final Report, (with redactions), 12th December 2013
<http://www.nbnco.com.au/content/dam/nbnco/documents/NBN-Co-Strategic-Review-Report.pdf>
2. 2010 NBN Implementation Study Report
<http://data.dbcde.gov.au/nbn/NBN-Implementation-Study-complete-report.pdf>
3. 2012 Analysys Mason report: “Review of the efficiency and prudence of NBN Co's fibre and wireless network design”, 2nd March 2012
<http://www.nbnco.com.au/assets/documents/a-c/analysys-mason-final-design-report-02-mar-2012.pdf>
4. Dr John Cioffi, Communications Day, 5th December 2013
5. NBN Co Network Design Rules, 18th September 2012
<http://www.nbnco.com.au/assets/documents/network-design-rules-2012.pdf>
6. Paul Spruyt and Dr Stefaan Vanhastel, “The numbers are in: Vectoring 2.0 makes G.fast faster”, Alcatel-Lucent Techzine, 4th July 2013
<http://www2.alcatel-lucent.com/techzine/the-numbers-are-in-vectoring-2-0-makes-g-fast-faster/>
7. CableLabs, DOCSIS 3.1, 29th October 2013
<http://www.cablelabs.com/specs/specification-search/?cat=docsis&scat=docsis-3-1>
8. Huawei, “Huawei Unveils the World's First G.fast over Coax Multi-user Access Prototype”, 14th October 2013 http://www.huawei.com/ilink/en/about-huawei/newsroom/press-release/HW_310334
9. Netflix CEO Reed Hastings interviewed by BBC's Rory Cellan-Jones, 9th January 2014
<http://www.bbc.co.uk/news/technology-25663472>
10. Hanhart, Rerabek, De Simone, and Ebrahimi, “Subjective quality evaluation of the upcoming HEVC video compression standard”, Proc. SPIE 8499, Applications of Digital Image Processing XXXV, 15th October 2012
http://infoscience.epfl.ch/record/180494/files/hanhart_SPIE2012_1.pdf
11. “H.265 : High Efficiency Video Coding”, ITU-T, April 2013 <http://www.itu.int/rec/T-REC-H.265>
12. McKell Institute, “Superfast Broadband: The future is in your hands”, October 2013
http://mckellinstitute.org.au/wp-content/uploads/2013/11/McKell_Broadband.pdf
13. NBN Co, “Test Agreement”, 4th December 2013
<http://nbnco.com.au/content/dam/nbnco/documents/sfaa-test-description-cell-site-access-trial-20131204.pdf>
14. CBS News, “Netflix app to stream 'Ultra HD' 4K video on new TVs”, 10th January 2014
<http://www.cbsnews.com/news/netflix-app-to-stream-ultra-hd-4k-video-on-new-tvs/>
15. Les Brown, “G.fast for FTTdp”, 22nd September 2012
http://www.itu.int/dms_pub/itu-t/oth/06/5B/T065B0000320009PPTE.ppt
16. TechWeek Europe, “BT To Trial Huawei G.fast FTTdp Copper Broadband Technology”, 21st October 2103 <http://www.techweekeurope.co.uk/news/huawei-gfast-fttdp-copper-bt-129855>
17. Tony Brown, Australian Financial Review, “UK G.Fast trial gives hope to Australian copper networks”, 29th October 2013
http://m.afr.com/p/technology/uk_fast_trial_gives_hope_to_australian_oki4JLC5zr3VTbozq8J6uJ

18. Lantiq, “Fibre-to-the-Distribution-Point (FTTdp)”, 2013
http://www.lantiq.com/uploads/media//2013_Lantiq_Aethra_FTTdp_Solution.pdf
19. Sckipio, G.fast Semiconductor Solutions www.sckipio.com
20. Smart Services CRC, iSee Video Collaboration
<http://www.smartservicescrc.com.au/products/market-ready/isee-video-collaboration>
21. Jorge Salinger, “DOCSIS 3.1 Plans and Strategies”, CED Webinar, 5th November 2013
<http://www.cedmagazine.com/webinars/2013/10/docsis-31-plans-strategies>

Glossary

4K TV	3840x2160 resolution, 4x pixels of HD TV
8K TV	7680x4360 resolution, 16x pixels of HD TV
ABHT	Average Busy Hour Throughput
ADSL	Asymmetric Digital Subscriber Line
ASIC	Application Specific Integrated Circuit
AVC	Advanced Video Coding
BSS	Business Support System
BT	British Telecom
CAPex	Capital Expenditure
CCAP	Converged Cable Access Platform
CIR	Committed Information Rate
CMTS	Cable Modem Termination System
CoS	Class of Service
CPE	Customer Premises Equipment
CVC	Connectivity Virtual Circuit
DFN	Distribution Fibre Network
DOCSIS	Data Over Cable Service Interface Specification
DSL	Digital Subscriber Line
DSLAM	Digital Subscriber Line Access Multiplexer
DVB	Digital Video Broadcasting
EPON	Ethernet Passive Optical Network
FDD	Frequency Domain Duplexing
FDH	Fibre Distribution Hub
FSA	Fibre Serving Area
FTTB	Fibre-to-the-Building/Fibre-to-the-Basement
FTTdp	Fibre-to-the-Distribution-Point
FTTLA	Fibre-to-the-Last-Active (Amplifier)
FTTN	Fibre-to-the-Node
FTTP	Fibre-to-the-Premises
Gbps	Giga bit per second

GPON	Gigabit Passive Optical Network
HetNet	Heterogeneous Network
HD TV	High Definition TV, 1920x1080
HEVC	High Efficiency Video Codec
HFC	Hybrid Fibre Coax
ISDN	Integrated Services Digital Network
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardisation Sector
kbps	kilo bits per second
LCD	Liquid Crystal Display
LED	Light Emitting Diode
LFN	Local Fibre Network
LNDN	Local Network Distribution Network
Mbps	Mega bits per second
MDF	Main Distribution Frame
MDU	Multiple Dwelling Unit
MHz	Mega Hertz
MIMO	Multiple Input Multiple Output
MPEG	Motion Picture Experts Group
OPex	Operational Expenditure
OSS	Operation Support System
OTT	Over-the-top
PABX	Private Automated Branch Exchange
PBX	Private Branch Exchange
PC	Personal Computer
PCM	Pulse Code Modulation
PIR	Peak Information Rate
PSNR	Peak Signal to Noise Ratio
PSTN	Public Switched Telephone Network
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RF	Radio Frequency
RFoG	Radio Frequency over Glass

RSP	Retail Service Provider
SID	Service IDentifier
SDU	Single-Dwelling Unit
SNR	Signal to Noise Ratio
SSID	Service Set Identifier
TC1	Traffic Class One
TCAM	Transmitter Controlled Adaptive Modulation
TDD	Time Domain Duplexing
UGS	Unsolicited Grant Service
UHD TV	Ultra High Definition, 4K or 8K
UNI-D	User Network Interface – Data
UNI-V	User Network Interface – Voice
VDSL	Very-high-bitrate Digital Subscriber Line
VoIP	Voice over Internet Protocol

Report Authors

Dr Craig Watkins

Dr Watkins has a Telecommunications Engineering R&D background forged at AT&T Bell Labs, Murray Hill, Motorola Corporate Research, Schaumburg, and Voxware Inc., Princeton. He has lectured in engineering communication systems at the ANU and has exposure to the Australian start-up environment through an electric vehicle enterprise, in addition to involvement in the US start-up industry through Voxware (an early VoIP and internet audio-streaming company). He holds several patents and has contributed to a number of international standards. In addition to networks and telecommunications, Dr Watkins maintains technical interest in, and a healthy passion for, the energy market, the education sector, establishment of innovation culture, effective research management, elementary physics, and the environment. Dr Watkins maintains broad consulting, strategy, and research involvements.

Ing. Kelvin Lillingstone-Hall

Kelvin is the CEO of OAK Telecom which started in 2000. During this time he has worked with a number of both private businesses and government organisations and helped them develop and implement long term broadband networking strategies. OAK Telecom resells the broadband networks of the major telcos in Australia through wholesaling intermediaries with their own data centres and comprehensive connections into each of the telcos. After graduating from RMIT in Communication Engineering as a Professional Engineer, Kelvin spent 30 years in Telstra working and gaining knowledge and experience on microwave system design, transmission planning, corporate planning, business planning, forward network planning and strategic planning as part of skilled teams which helped develop and implement Telstra's extensive networks around Australia. Early in his Telstra career Kelvin went on a two-year engineering scholarship to GEC Telecommunications in Coventry, England. Subsequently he spent six months in Indonesia with the Foreign Affairs Department working on the design of the Trans Sumatra Microwave System. Kelvin later moved from Telstra to gain sales and marketing experience that lead him to work as a professional consultant to a number of medium to large corporations helping them to develop their voice and broadband networking arrangements. Kelvin has been deeply involved over many years in the Victorian and national committees of Engineers Australia and Professionals Australia (formerly APESMA). Kelvin is passionate about providing excellent service to his business clients as well as helping to promote Professional Engineers and their work in Australia.

Reviewers

Peter Hitchiner

Peter Hitchiner has forty years of experience in the telecommunications industry including roles as Director and State Manager for six years with a leading telecommunications consulting business and as General Manager for four years with a telecommunications project management business. That engagement included implementing a major optical fibre cable system in North America (\$US250M) and planning of a major fibre optic infrastructure investment Australian Fibre Networks (AFN). Peter has extensive knowledge of optical fibre systems from a technical as well as implementation perspective. Peter worked extensively with legal and financial advisors in particular in the development of the Australian Fibre Networks proposal. Of particular relevance to NBN he undertook (in early 2009) a scoping study and strategy for fibre to the premises in a major residential/industrial development in SE Queensland involving up to 50,000 new households. Peter has an occasional advisory role with Smart Services Cooperative Research Centre and this provides important insight into applications and services which are supported by broadband infrastructure. Peter is a Fellow of the Institution of Engineers, Australia (EA) and immediate past Chair of the Board of the Information, Telecommunications and Electronics Engineering College and is a Past President of the Sydney Division. Peter represented Engineers Australia on the Fibre in Greenfield Stakeholder Reference Group considering matters related to the implementation of Fibre to the Premises in greenfield developments. Peter was a member of the steering committee for a project which resulted in the publication (December 2007) of EA's Telecommunications Infrastructure Report Card which assessed the state of broadband and mobile telecommunications infrastructure through Australia.

Dermot Cox

Dermot Cox is an independent commentator on the Australian NBN with a special focus on designing broadband policy which underscores service outcomes and leverages existing super-fast broadband assets such as cable broadband (HFC). He has submitted papers to the Senate Select Committee on the NBN and appeared before the Select Committee on two previous occasions. Professionally, he advises firms on projects to improve sales effectiveness such as improving bid marketing, undertaking sales audits, and the application of sales methodology technology to improve performance. Previously he joined C-COR Broadband, an Australian systems integrator for HFC cable broadband, in 2006. He was intimately involved in engaging senior carrier and regulatory executives on advances in cable technology. In this role he led the firm's advocacy and strategic marketing across the region. Before C-COR, he held key international sales and market development leadership positions with Nortel primarily focussed on wireless and wireline service providers. He developed business cases for new market introductions for leading-edge technologies spanning high-capacity optical transport, network telephony, mobile 3G, next generation IP architecture networks, and terabit IP transport. He is also on the national committee of management for The Wilderness Society. He received his Masters degree in marketing from the Monash University and is an Associate Fellow of the Australian Marketing Institute and a Certified Practising Marketer.