Robert Kenny & Tom Broughton

Domestic bandwidth requirements in Australia
A forecast for the period 2013-2023

26 May 2014

COMMUNICATIONS CHAMBERS

About the authors

Robert Kenny advises companies, regulators and policy makers on issues of TMT strategy and policy. He is a regular speaker on these topics and the author of numerous papers (academic and professional). Before co-founding Communications Chambers he was MD of Human Capital, a consulting firm. Prior roles include heading strategy and/or M&A for Hongkong Telecom, Reach and Level 3 (all multi-billion dollar telcos). He was also a founder of IncubASIA, a Hong Kong based venture capital firm investing in online businesses.

Tom Broughton has worked as a strategy consultant to the telecoms and media industries for over a decade. Prior to becoming an Associate with Communication Chambers he worked at Ingenious Media, Human Capital and Deloitte. Tom is a chartered accountant and holds a degree in Econometrics.

Acknowledgements & Disclaimer

This report was commissioned by The CIE on behalf of the Vertigan Panel. It has greatly benefited from review and proprietary data provided by a range of industry participants, and we gratefully acknowledge their support.

However, the opinions offered herein are purely those of the authors. They do not necessarily represent the views of The CIE, the Panel or any industry participant, nor do they represent a corporate opinion of Communications Chambers.

Contents

1. Executive Summary

This paper sets out the methodology and results of a model that seeks to forecast Australian domestic requirements for broadband capacity. By this we mean the capacity likely to be used by consumers – this is distinct from the bandwidth consumers might be willing to pay for (which could be more or less).

The model is technology agnostic – it is simply focussed on demand, and does not consider how that demand might be met (for instance, via a range of fixed technologies or by wireless).

The model is also narrowly focused on access-network bandwidth. There are numerous other factors, ranging from in-home wifi capacity through to web page design which can have an impact on the ultimate user experience. Thus the access bandwidths discussed in this report are *necessary* for a good user experience, but they are not *sufficient*, in that other factors could degrade this.

Objectives

Given the worldwide interest in increasing broadband speeds, it is perhaps surprising that there is so little methodical work[1] on what households' requirements might be in this area. In developing our own view, we have felt that the following are vital:

- A rigorous approach to determining the speeds required by individual applications, in particular by investigating the speeds recommended by the leading providers of the service in question.
- Accounting for changes in the required speeds over time. In many cases, requirements will rise – for instance, consumer expectations of acceptable download times will likely shorten. However, requirements may also fall. In particular, constantly improving video compression means that (for a given video definition) required bandwidths will decline.
- Reflecting variation in demand across households, particularly that driven by household size. Approximately 58% of Australian households contain only one or two people.[2] The average usage of such households will inevitably be lower than that of larger households.
- Building a quantified view of the probability and likely duration of 'app stacks'. Much discussion of future

[1] See however Analysys Mason (for BT), *International benchmark of superfast broadband*, 27 November 2013, which, at a higher level, uses some of the same approaches we deploy in our forecast

[2] Communications Chambers analysis of ABS data

bandwidth needs has been anecdotal, such as "imagine a household doing X, Y and Z simultaneously". While any such 'app stack' is conceivable, that does not necessarily mean it is likely or regular.

- Providing a picture of the duration of the highest levels of demand in a household. This enables trade-offs to be made. For instance, if a household only requires a high level of bandwidth for a short period, will they (or society) feel it is worth the incremental cost to secure that high bandwidth, or will they instead choose to tolerate a period of degraded quality in exchange for a lower cost?

Approach

To meet these requirements, the model takes the following approach. It is bottom up, beginning with a set of 13 different categories of applications that cover the vast majority of ways in which the domestic internet is used. Some, such as web use, are broad, covering everything from Facebook through to online taxes to certain varieties of cloud services such as salesforce.com.

Figure 1: Modelling approach

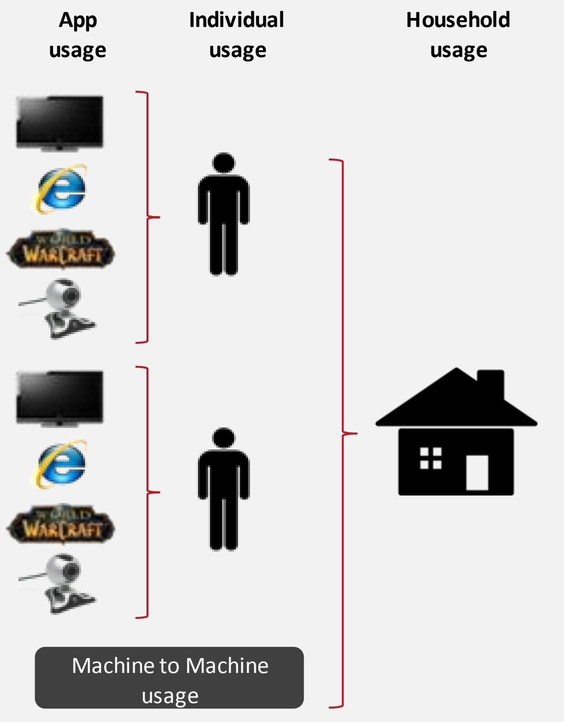

Based on best available data, we have made assumptions regarding both the bandwidth needs and the volume of usage of each of these applications. Note that we have not necessarily based these on today's current usage. In some cases we believe that current usage was constrained by current bandwidth, rather than reflecting what might be reasonably expected absent this constraint. To take one example, while downloading a console game might typically require a user to wait overnight or longer, we have set expectations to 60 minutes, with this figure decreasing over time.

However, in making such assumptions, we have had to take a middle line between today's expectations and a 'perfect world'. In a perfect world, everyone would be able to download everything in seconds. However, this would imply that everyone 'needs' gigabit speeds today – we do not feel this is a meaningful or helpful approach. (As a practical matter, other elements of the chain from the content provider's server to the consumer's device might well not be able to provide gigabit speeds, even if the access network could).

We have combined our usage profiles of the various applications into usage profiles of individuals. In doing so, we have taken a probabilistic approach to combining apps. For example, if app A is being used for 50% of the time, and app B is used for 40%, then we would expect the two to be used simultaneously for 20% of the time.

We then further combine these individual profiles into household profiles (again taking a probabilistic approach), depending on household composition. This is based on 16 household types (for example, single adult, two adults, two adults one child and so on), and also on the type of primary TV set (SD/HD/4K - extremely high resolution "4K" TVs have the potential to be an important contributor to bandwidth demand). We also disaggregate households out into high, medium[3] and low categories based on their propensity to use the internet. In combination, these splits build to a total of 192 household types.

Figure 2: Household types

16 *Demographic types*	X	*4* *Usage intensities*	X	*3* *Main TV types*	=	*192* *HH types*
1 adult, 2 adults, 1 adult, 1 child …		High/Med with software DL & video UL/Med/Low		SD/HD/4K		

These 12 profiles are then combined into a picture of the national mix of demand. Note that we believe these profiles (and the underlying usage assumptions) cover the vast majority of cases. However, they will not address *all* possibilities – for instance, people working at home on professional animation or analysing astronomical data sets might need to up- or download very large files, which we have not covered. Such cases are outside our scope.

In undertaking the above analysis, we have focused on the 'busy hours'. Bandwidth demand is obviously driven by peaks, not the average speed required, and we have therefore focused on the busiest four hours per day (in many households this is likely to be 7-11pm or later). We have assumed that 50% of usage occurs in these four hours – put another way, this assumption means that the rate of usage in this period is five times the rate of usage in the rest of the day, and thus tall 'app stacks' are much more likely.

[3] The medium category is then further split into those who do and do not participate in large software downloads and video uploads

Results

Figure 3 shows the usage profiles generated by the model for three sample household types in 2023. For the single person low usage household without HD TV, the connection is in fact idle for much of the time, but has several hours at 5 Mbps and short periods in the range of 8-10 Mbps. Conversely, the four adult house is seeing appreciable usage almost constantly during the busy hours, at times of 20 Mbps or even higher.

Figure 3: 2023 usage profiles, three sample household types

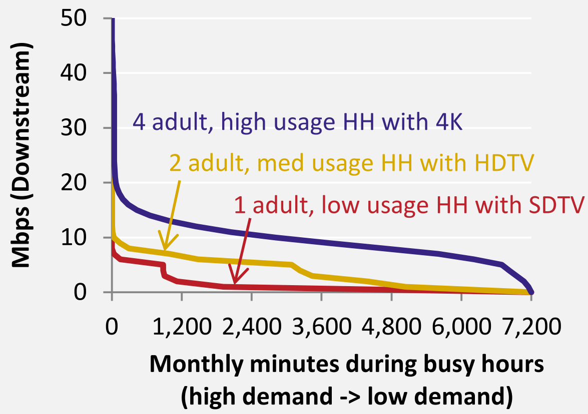

We have used such profiles of household usage to determine '4 minutes excluded monthly' bandwidth demand. By this we mean the bandwidth that would be necessary to serve all but the four busiest minutes in the month (or one per week for each household) – that is, the four minutes at the extreme left of Figure 3. We believe this metric is useful, since it seems plausible that consumers would accept one minute per week of degraded performance (such as a video stream with briefly lower resolution) if it brought them any cost saving.[4]

On this '4 minutes excluded monthly' basis, bandwidth demand grows as shown in Figure 4. As of 2023, the median household requires bandwidth of 15 Mbps (the arrow on the chart), while the top 5% have demand of 43 Mbps or more.

Figure 4: Downstream Bandwidth demand distribution (4 mins excluded basis)

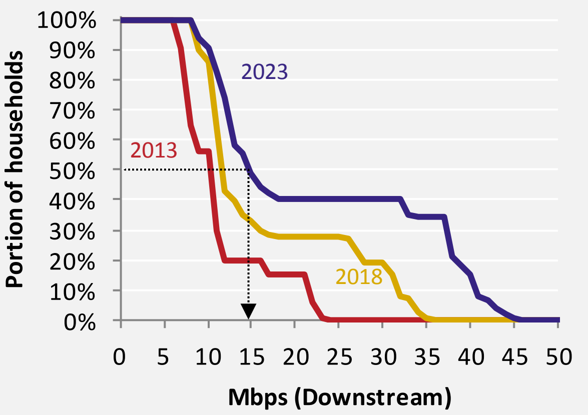

Note that in offering figures for median demand, we are not suggesting that access network capacity should be based on this metric – by definition, such a network would (to some extent) disappoint 50% of households. Rather, access capacity should be driven by higher end users. Whether this means the top 1% or the top 5% (or some other figure) is a matter of judgement. However, the median demand figure illuminates the

[4] The choice of four minutes is of course arbitrary – we provide sensitivities to explore the impact of varying this assumption. We note that some commentators have taken the approach that all demand, no matter how brief, must be met. However, this not consistent with how the rest of the telecoms network is provisioned, nor indeed how other forms of infrastructure such as roads are provided.

extent to which building a higher capacity network to benefit these high-end users involves a cross-subsidy from more typical users.[5]

The figures of 15 Mbps and 43 Mbps for median and top 5% demand may seem low, particularly by comparison to the results of more informal work in this area. However, the most common type of household comprises just two people. Even if those two are each watching their own HDTV stream, each surfing the web and each having a video call all simultaneously, then (in part thanks to the impact of improving video compression) the total bandwidth for this somewhat extreme use case for that household is just over 14 Mbps.

Interpreting the results

We do not offer our results as 'the truth'. There is appreciable uncertainty, and these figures are at best a mid-point prediction. As with any model, ours depends on a wide range of assumptions (set out in this paper), some of which are unavoidably arbitrary. Others may quite legitimately take a different view.

One prime example is the set of assumptions about content download expectations. For example, we have assumed that movie downloads happen at least four times faster than real time[6] in 2023 (so 30 minutes for a 2 hour movie), but one might instead assume users will require that movies download in 1 minute.

Specific assumptions aside, the key themes of our approach that contribute to our conclusion are:[7]

- A rigorous analysis of the probability of app stacks
- Reference to the actual bandwidth requirements of individual apps (and their increase or decrease over time) rather than loose estimates
- An understanding of actual household demographics rather than on a notional 'typical' household (which is frequently anything but)
- An understanding of the duration of peak demands

Of course our modelling exercise is only a contribution to a meaningful discussion about Australian bandwidth needs, rather than its conclusion.

---

[5] This cross subsidy is of course eliminated if these high end users pay a premium such that they fully cover the incremental cost of bandwidth which only they use

[6] Note that this is a worst case assumption, for download times in the busiest minutes of the month – in practice, most movies would download far faster

[7] See page 56 for a more detailed discussion

2. Introduction

Broadband has significant social and economic benefits – its capabilities are vastly greater than dial-up. Basic broadband is now widespread (though with some variability in available speeds), and attention has long since turned to 'superfast' broadband and issues of uptake and usage.

There are several ways to deliver superfast – via fibre to the home (FTTH) or the cabinet (FTTC), via cable broadband and via wireless networks (amongst others). These approaches have their pros and cons. FTTH is more expensive, for example, but can deliver higher speeds, lower latency and so on.

Critical to assessing the cost/benefit trade-offs of various superfast technologies is an understanding of likely future speed requirements. For instance, if FTTC is chosen over FTTH, how many households are likely to have their usage constrained by FTTC's lesser speed (and what particular applications are likely to be most affected)?

However there is a surprising lack of forecasts in the public domain on the expected demand for bandwidth. There are certainly *traffic* forecasts. Cisco's Visual Networking Index forecast is widely cited, for instance. However, while valuable, these tell us little about the need for bandwidth, particularly at the edge of the network (the last mile to the consumer).

This is because the edge network has (on an average basis) very low utilisation – see Figure 5. Even during busy evening hours the average connection in Australia is only used to 1.7% of its capacity. The existing access network could potentially absorb great amounts of additional traffic without needing any upgrade at all, and thus traffic forecasts tell us relatively little about bandwidth requirements.

Figure 5: Busy hours utilisation of Australian broadband

Data per connection per month[8]	47.4 GB
Of which in busy hours (6-12pm)[9]	30%
Data per connection in busy hours	14.1 GB
Traffic per hour in busy hours	80.21 MB
Average usage	0.18 Mbps
Average ADSL line speed[10]	12.7 Mbps
Average utilisation[11]	1.7%

However, though the *average* utilisation is only 1.7%, this certainly does not mean that

[8] Communications Chambers analysis of data from ABS, *Internet Activity Australia, December 2013*, 7 April 2014
[9] Estimate, based in part on UK data – see Ofcom, *Infrastructure Report 2013 Update*, 24 October 2013
[10] Department of Communications, *MyBroadband data cube*, 11 March 2014. For reasons of data availability, higher speeds available on HFC and via the NBN are excluded, so this figure is an underestimate of overall Australian broadband line speeds
[11] The calculated utilisation also allows for the fact that available bandwidth may be 13-15% lower than sync speed

there is no need for additional bandwidth. This is because per-home bandwidth utilisation is far from smooth – it varies dramatically over short periods of time. To take a simple example, on Monday the members of a particular household all may be out, with a home bandwidth need of 0 Mbps. On Tuesday evening they may be in, watching simultaneous catch-up TV streams, requiring 5 Mbps. It is of course this latter peak that sets their perceived bandwidth need, not the average of the two days.

A given household can have variable bandwidth needs over short periods of time, but there is also significant variance between households. The retiree living alone will likely require far less bandwidth than a family with two parents and three teenagers.

Finally, there is of course variance over the medium and long term. Internet usage rises year-by-year, with (for example) levels of video consumption today far above those of five years ago.

To understand bandwidth requirements it is therefore not sufficient simply to look at averages – we need to consider how many households will need a given bandwidth, how often. For instance, one household might need 10 Mbps for one minute per month, another might need it for one hour per day. The former might choose a lower bandwidth product, the latter would likely purchase at least 10 Mbps. This paper (and the associated model) seeks to provide forecasts of such variation.

Both the paper and the model draw extensively on work Communications Chambers undertook last year for the UK's Broadband Stakeholder Group, an advisor to the UK government funded by a wide range of interested parties, including the country's Department of Communications, Media & Sport; leading telcos and ISPs; mobile operators; broadcasters; and equipment manufacturers. Several of these bodies provided peer-review and proprietary data for our UK forecasts, and this report benefits indirectly from that support.

In preparing this report we have made numerous additions and amendments to the UK model, including:

- Changing input assumptions to reflect Australian data (wherever possible)
- Updating the model for international market developments (such as lower than expected required bandwidth for 4K TV and higher than anticipated file sizes for console games)

- Making revisions in light of feedback received after publication of the UK model
- Adding additional detail on areas of particular relevance to the Panel's consideration
- Making amendments in light of feedback and data received from Australian market participants

We are grateful to the Australian RSPs and content players who provided commentary on the methodology and/or proprietary data which allowed us to sanity check our inputs and outputs.

As with any forecast, this one is certainly open to debate and alternate assumptions. We have sought to provide the sources for our underlying assumptions, or where they were simply a matter of judgement to explain our rationale. We do not intend the outputs to be 'the truth', but rather simply a credible middle case (and we provide sensitivities to this middle case).

Indeed, the value of models is more often in the rigorous analysis of a problem that they require, rather than in the results themselves, and we have therefore also set out our methodology. This too is, of course, open to debate.

The model is also limited in scope – it considers domestic demand (not commercial); looks only at bandwidth not other aspects of technical requirement; and by definition only forecasts types of demand we know about today (though it certainly includes applications we do not yet know about, if they would fall under one of the types of demand such as 'web surfing').

A note on terminology: we use 'conservative' to mean 'unlikely to lead to too low a bandwidth forecast' - that is, an assumption is conservative if it leads to higher bandwidth than the alternative.

3. Macro growth trends

Very broadly, the peak bandwidth requirements of an online household are driven by:

- The number of people using the internet in that household
- The length of time that each user spends using the internet, and the degree to which they multitask (since this drives the level of overlapping usage, for individuals and households)
- The bandwidth of the applications used while online

The last of these continues to increase, as the mix of usage shifts increasingly to video.[12] However, the first is now essentially flat, and the second may be approaching saturation in the years ahead.

Internet users per online household

In recent years the number of internet users in Australia has continued to grow, from 11.2m in 2007 to 15.4m in 2013. The portion of households online has grown from 64% to 83% over the same period. [14]

However, recent growth in online households has disproportionately come from older, single person households. Consequently the average number of internet users per online household has fallen slightly, stabilising at around 2.1.[15]

Figure 6: Internet users per online household[13]

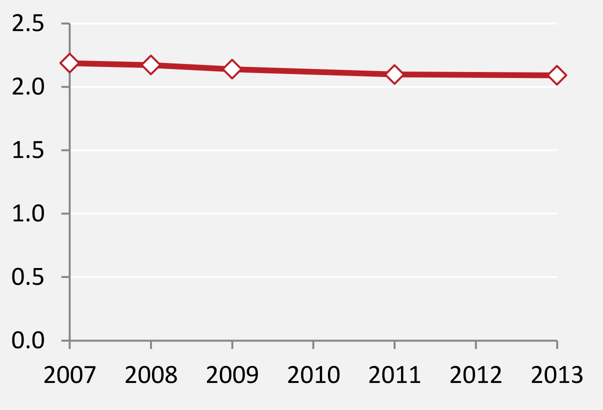

[12] Though on a standalone basis video may be approaching 'peak bandwidth', given ever improving compression and the limits of the human eye
[13] ABS, *Household use of Information Technology, Australia 2012-13*, 25 February 2014
[14] Ibid
[15] Ibid

Time spent using the internet

There are obviously natural limits to how much time an individual can spend using the internet. According to Nielsen the average online Australian currently uses the internet for 39 hours per month (or 1.3 hours per day) via laptops and desktops, but it may be that this form of usage is approaching saturation (see Figure 7). It is only slightly behind North American levels of usage (42.8) and well ahead of the average in Europe (26.9).[17]

Figure 7: Monthly internet usage via desktop & laptop, hours/user[16]

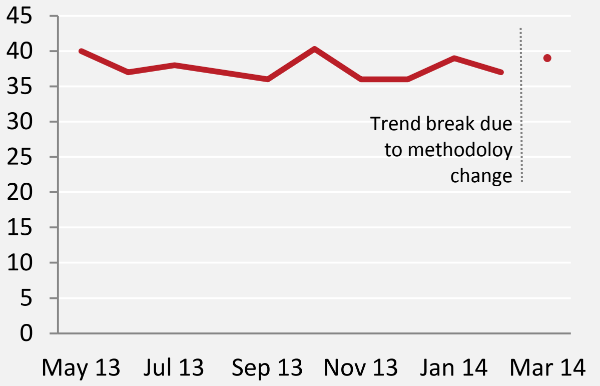

That said, while time online via laptops and desktops is flat to falling in developed markets, growth is being driven by increasing consumption via smartphones and tablets.[18] For instance, average video consumption on mobile devices is now 3 hours 43 minutes, up from under half an hour three years ago.[19]

Resulting traffic

As we have noted, by itself traffic growth is only loosely related to bandwidth requirements over time. Nonetheless, it is helpful to look at historic growth in traffic per line (see Figure 8).

Figure 8: Traffic per fixed BB line, GB/month[20]

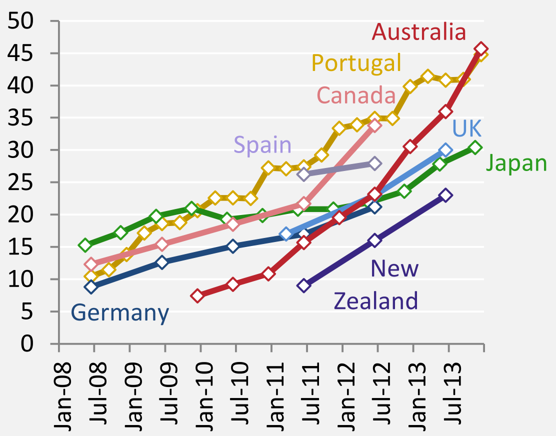

Clearly there is substantial variation between countries. For example, Japan (despite widespread FTTH) has relatively low traffic. While once ahead of the comparator countries, it has now fallen well back, perhaps because of heavy usage of mobile data networks. Conversely Hong Kong (which also has widespread FTTH) is omitted from the chart since its traffic is so high, at 100 GB for December 2013.[21]

---

[16] Nielsen, *The Australian Online Landscape Review, February 2014*, 26 March 2014 (and respective prior months). Note that due to a change in Nielsen's methodology, longer run data is not available
[17] Comscore, *UK Digital Future in Focus*, 14 February 2013
[18] Ofcom *International Communications Market Report 2013*, 12 December 2013
[19] OZTAM / Nielsen, *Australian Multi-screen Report – Q4 2013*, 11 April 2014
[20] Source from relevant national regulatory authorities or government statistical services. Figures are average for both business and residential lines, except for the UK which is residential only. Australia traffic is for download only – upload also included for other countries
[21] OFCA, *Statistics on Customers of Internet Service Providers ("ISPs") in Hong Kong*, February 2014; OFCA, *Customer Access via Broadband Networks*, 30 April 2014

Australia has seen strong growth, particularly since early 2011, perhaps as a result of the 90% reduction in cost per GB in 2009 and 2010.[22] This appears to have triggered a material change in behaviour, including a rapid rise in online video consumption, which increased from 2:42 hours per month per person in Q1 2011 to 9:35 hours in Q4 2013 (equivalent to a 59% annual growth rate).[23] Other forms of traffic will also have seen rapid growth – for instance, we estimate current per-line Bittorrent traffic is equivalent to *total* per line traffic of two years ago.

As a consequence Australia's usage is now the highest in the comparator set (Hong Kong aside). Moreover, the chart understates Australia's relative traffic since the ABS figures on which it is based cover only downstream traffic – for other countries the traffic shown is both up- and downstream. (Upstream traffic might typically be 15 to 25% of total).

All countries are seeing growth in per-line traffic, but broadly speaking that growth appears linear rather than exponential.

With this context in mind, we now turn to our methodology for forecasting future bandwidth requirements.

---

[22] John de Ridder, *Australian retail broadband – price competition has stalled*, 30 September, 2013

[23] OZTAM / Nielsen, *Australian Multi-screen Report – Q4 2013*, 11 April 2014; OZTAM / Nielsen, Australian *Multi-screen Report – Q4 2011*, 20 February 2012

4. Modelling approach and structure

Introduction

To estimate an individual household's usage we have taken a 'bottom up' modelling approach - starting with a set of applications used by individuals, combining these to build a profile of individual usage and then combining these individual profiles to get a picture of household usage.

Figure 9: Modelling approach

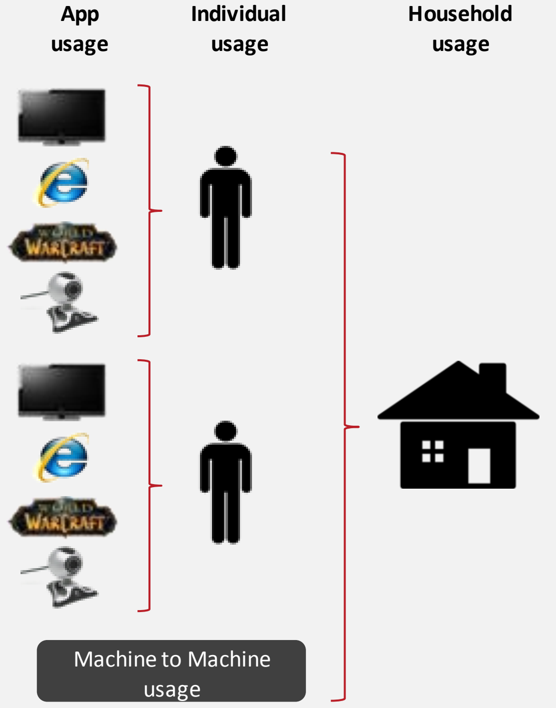

Scope of the model

The focus of the model is in-home usage, and specifically the peak busy hours of that usage. We have not addressed internet use either on the macro-cellular network (or wifi) while out of the home, nor have we considered workplace usage (home-working aside). While focused on in-home usage, the model is technology agnostic. We forecast the access bandwidth requirement, not how that demand might be met (for instance, via a range of fixed technologies or by wireless). We also do not consider bandwidth for backhaul (connecting the access network with the wider internet). The cost of such capacity, such as NBN CVC cost, is material, and this may therefore in practice represent a constraint on usage, even if the access network has sufficient bandwidth.

The model also does not explicitly treat devices – our focus is on the usage (for instance, web surfing) rather than the device on which that usage happens. We believe that this approach is preferable, since it is the person that is likely to be the constraint in the system. The number of internet-capable devices may carry on rising, but as a practical matter a person is only going to be able to use a certain number simultaneously. That said, we have built in an ongoing increase in web usage (for example) in part because increasing adoption of devices will enable more frequent internet-usage around the home. Tablets are important in making video available beyond the living room, for example. Moreover, mobile devices are a key enabler of the combination of applications (such as web usage plus Internet TV consumption) that the model analyses.

Note that while we believe the model addresses the vast majority of households, it does omit a small number of rare or extreme cases – households with more than four adults or someone running an animation studio from home, for example.

Applications as building blocks

The basic building block of the model is the application. We set out here our broad approach to applications. In section 5 we discuss our thinking and assumptions for each application type in detail.

The model treats 13 different application types explicitly. These include items such as web use, high definition video calls, 4K TV and so on. We have treated explicitly those applications that would have the biggest impact on bandwidth demand. Some activities requiring little bandwidth, such as internet radio and e-metering have been grouped under 'low bandwidth' applications.

Note that the applications in the model are broad in the range of activities each covers. For instance, web surfing covers everything from use of salesforce.com through Facebook through to humanservices.gov.au. Our focus has been on the technical characteristics of the traffic which drive bandwidth (for instance, page load times for web traffic), rather than on the precise use to which (say) web pages are put. This has the great virtue of obviating the need to predict the next Facebook.

One consequence of this approach is that we do not explicitly forecast bandwidth for groupings of applications such as e-health. Rather, such uses are embedded within other applications in the model. For instance, medical telemetry is included within low bandwidth apps, a video consult with a doctor within HD video calls and so on.

Categories of applications

We have put our applications into four categories: primary, secondary, web use and low bandwidth.

Figure 10: Illustrative individual usage over a day

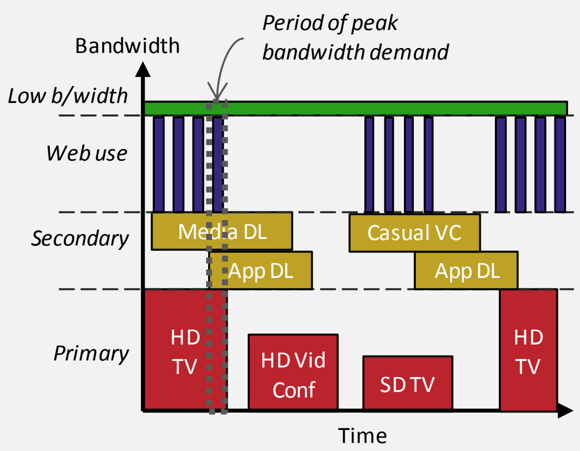

Primary applications are those apps that are primarily used 'one at a time' by a given individual, in particular Internet TV[24] consumption (though they will be used in parallel with non-primary apps).

Secondary apps are amenable to multitasking – for instance, launching a movie download and then continuing with other activities, such as web use. Note that to be conservative (in terms of not underestimating demand) we have included within 'secondary' applications uses such as video calling that may generally be a primary activity but which, occasionally, might be used in parallel with (say) Internet TV consumption.

Web use is in its own category because of its bursty nature as individual pages are loaded – this means more bandwidth is required than average traffic would suggest. Our fourth category is the catch-all 'low bandwidth'.

The full set of applications treated is as follows:

Figure 11: Application types

Primary	Secondary	Web	Low bandwidth
Internet TV (SD, HD and 4K)	Cloud storage	[Surfing, excluding use of video sites such as YouTube]	[Covers a range of applications not explicitly treated, including e-metering and other machine-to-machine, online radio, online gaming etc]
HD video calls	Content downloads		
YouTube etc.	P2P / BitTorrent		
Streamed gaming[25] / HD interactive	Mobile OS downloads		
	Software downloads		
	Non-HD video calls		
	Content uploads		

[24] We use Internet TV to mean streamed professional content, regardless of whether delivered to a TV set or a mobile device. In the Australian context today, this primarily means catch-up services such as ABC's iview, or on-demand services such as Quickflix or Netflix. In some markets broadband is also used as a distribution channel for linear (scheduled) TV, and this may happen in time in Australia also

[25] Streamed gaming refers video streamed from a central server, dynamically generated in response to player actions. Online gaming refers to multiplayer games where the internet is used to share data on the game state, not video. The latter requires low latency, but not high bandwidth. Even a relatively demanding game such as *Halo 3* only requires 60 kbps. Andreas Petlund et al, *Network Traffic from Anarchy Online: Analysis, Statistics and Applications*, 22 February 2012

This set of applications includes (but is not limited to) all the key drivers of traffic in the network busy hour today.[26] It also includes applications such as streamed gaming, HD video calls and 4K TV, which are not important today but which are often cited as drivers of future demand. Other uses that are sometimes cited as future drivers are home-working and e-health. As noted above, while these are not explicitly included, their components (HD video calls, cloud storage, telemetry and so on) *are* included in our application set.

Focus on the busy hours

In forecasting bandwidth requirements, the model focuses on the busy hours (those when the internet is being used most intensively), since these will be when peak demand occurs. The model begins with monthly usage per person of the various applications, but then assumes that for each application, 50% of usage will take place within the four busiest hours each day for the household.

Based on industry data, we believe approximately one quarter of network-wide traffic flows in the four busiest hours. Our figures imply much greater concentration, but that is appropriate both because traffic is likely to be more concentrated at the household level than it is at the network level,[27] and because the network wide busy period is extended by different local peaks in different time zones.

By concentrating the traffic, we greatly increase the likelihood of overlapping usage (both for an individual and across individuals), thereby upping the peak bandwidth requirement.

Note that the model is agnostic as to *which* four hours in the day are the busiest – for a household with someone working from home, it might be during the day; for another it might be in the evening as the family settles in to watch Internet TV.

Combining usage

Having understood the levels of usage in the busy hours of the various applications, the next step is to combine these usages into a picture of total demand, based on the extent to which they overlap.

---

[26] Based on Sandvine, *Global Internet Phenomena Snapshot: 2H 2013*, 11 November 2013

[27] This is because not all households have the same busy hour. Consider four households, one using the internet only from 7-8pm, two using it only from 8-9pm and one using it only from 9-10pm. From a network perspective the busy hour is 8-9pm, and this contains 50% of total traffic. However, for each individual household, 100% of its traffic is in its respective busy hour

A household in which one person watches an HDTV stream (at 3 Mbps) in the afternoon and another watches HDTV in the evening has a peak demand of 3 Mbps. Conversely, if the two individuals happen to watch their respective programming simultaneously, then the peak demand is 6 Mbps.[28] Thus an understanding of the expected likelihood, frequency and duration of overlaps is key.

To build a profile of combined usage, we have taken a probabilistic approach. For example, if secondary app A is used for 30% of the time during the busy hours, and secondary app B is used for 40%, then they will (on average) be used simultaneously for 12% of the time.

This approach is used both to combine app usage to build individual user profiles, and to combine user profiles to build household profiles. In practice, this is a complex set of calculations that ultimately embeds all possible combinations of applications, the total bandwidth for each combination and the expected duration.

Note that there are reasons to believe this probabilistic approach may potentially understate or overstate overlapping usage. It may understate it to the extent to which certain applications have a particular 'affinity' for each other. For instance, a user may potentially be *more* likely to surf the web while watching TV than they would be to surf at other times. (As a practical matter, we note that the model anyway results in substantial time for this particular combination, given the high volumes of both TV and surfing). Another possibility is that Internet TV streams will tend to overlap, since user A watching Internet TV may encourage user B (with different tastes) to seek out their own stream to watch.

Conversely, the probabilistic approach may overstate overlaps if applications 'repel' together. For example, if person A in a house is watching an Internet TV stream, it may be less likely that person B's own Internet TV usage is simultaneous, since they may instead be watching the first stream with person A. (Or person A may be watching Internet TV precisely because person B is watching broadcast TV – without using bandwidth - on the main set. Broadcast is still by an enormous margin the predominant mode of TV consumption, and seems likely to remain so).

[28] Note that the model makes the simplifying assumption that bandwidth usage is additive – one stream of X Mbps and another of Y Mbps will together need X + Y Mbps. In practice, two streams can sometimes interact in subtle and unexpected ways (both in the access network and more generally), meaning that even with ample bandwidth performance can be unsatisfactory. See the papers of Predictable Network Solutions for a more detailed discussion

Applications may also be more likely to 'repel' each other as they 'stack'. This is because of the human limits of multi-tasking. While it is perfectly feasible to (say) surf and watch Internet TV at the same time, each additional task becomes harder. For instance, if a user then receives a video call, it is likely (though not certain) that they would either pause the TV or their surfing. In other words, this triple multitask is less likely than simple probability (and the model) would suggest. This potential overstatement in the model is particularly important, since it is these app stacks that drive peak bandwidth demand in the busiest minutes of the month. This is one way in which the model is conservative (in the sense of unlikely to understate demand).

Variations in usage

The model does the above calculations not just for a typical household, but for a wide range of household types. Understanding variation in usage between different households is vital - to take a simple example, a level of bandwidth that is sufficient for a single person household may not be enough for a household with two adults and three children.

We consider variation by household demographics (number of adults and children – 16 types of household); whether they use SD, HD or 4K TV (3 types) and whether the individuals in the household are light, medium[29] or heavy users of the internet (4 types). Since these categories multiply up, the model considers 192 different household types overall, enabling us to understand not just what the typical household requires, but also to understand the spread from the most demanding to the least demanding household type.

Demographic types

Based on ABS statistics,[30] we have estimated the mix of household types as shown in Figure 12. Note the importance of one and two person households which represent 58% of the total.

Figure 12: Modelled household demographics

		Children			
		0	1	2	3+
Adults	1	24.3%	2.5%	1.9%	1.0%
	2	31.3%	7.6%	10.5%	5.5%
	3	7.1%	2.1%	1.1%	0.5%
	4	3.0%	0.8%	0.4%	0.2%

That said, this is the mix of all households, not those online. The elderly will be heavily represented in the one and two adults only households, and they are less likely to be online today. We have not modelled the impact of this, on the basis that by 2023 internet adoption can be expected to be very high across all demographics. For similar reasons, we

[29] This medium category is further split between those who do and do not upload videos and download large software files

[30] ABS, *Table Builder*. Note that some estimates have been required for multi-family households

have assumed all individuals in a household are online. Even today, this is likely true of most online households.

Turning to usage, in considering how the average usage of a child compares to that of an adult, we have taken the following into account:

- 'Child' here refers to those under 15 years of age; or family members aged 15 to 24 years in full-time education. Of this category, 89% are aged 16 or less, and 22% are aged three or less. This latter group of the very young will clearly pull average usage down.[31]
- Our focus is on the four busy hours, which in many households are likely to be between 7 and 11 pm. Children's usage may be somewhat less likely to occur in this period than that of adults.
- According to Nielsen, average online time for those aged 2-17 is 10 hours per month compared to an overall average of 37 hours.[32] However this group's video consumption, at 6.4 hours, is closer to the overall average video consumption of 7.7 hours.

In light of the above, we have assumed that average application usage by a child in the busy hours is half that of an adult. Note that the impact of children's usage on overall results is anyway reduced since they are only 23% of the population, and are only present in 34% of households.

Low, medium and high users

Even households with identical demographics may have different levels of usage. Such variation may be driven by levels of technical sophistication, habit, availability of devices and so on. To reflect this variation we have built into the model a low/medium/high spread of usage at the household level (on top of any variation caused by household composition). We have arbitrarily split households 40/40/20 across these categories, and assume low users have half average usage, and high users have double.

TV type

Our final form of household variation is related to the household's primary TV type, which we split between standard definition, high definition and 4K. 4K is of particular significance here – consumption of 4K Internet TV can create a high bandwidth need,

[31] ABS, *Table Builder*

[32] Nielsen, *The Australian online landscape review – February 2014*, 26 March 2014

but the take-up of such devices across most households is uncertain.[33]

Development over time

The model develops a ten year forecast for unconstrained[34] bandwidth requirements. Changes in requirements are driven by a range of factors, including:

- Increasing time spent using applications
- Varying bandwidth required by individual applications (for instance, rising for web use as web content includes richer images and more video, falling for SD TV as video compression improves)
- More widespread use of certain applications (for instance, uptake of 4K TV online)
- Rising user expectations (for instance, increasing impatience with time taken for software downloads)

Conclusion

While the modelling of the combination of applications to build household usage profiles is complex, it is to some extent mechanistic. This aspect of the model is not in itself highly assumption-dependent, but rather depends on the mathematics of overlapping usage.

However, it builds absolutely on the assumptions regarding the usage characteristics of individual applications. We now turn to this vital area.

[33] Note that other devices besides TVs (such as tablets) may have a display that is notionally 4K resolution. However, we set these aside since without a large screen seen close up, it is in practice impossible to detect the higher resolution, and by extension bandwidth for streaming below 4K levels would have no detectable impact on the consumer

[34] By this we mean that we have not scaled back bandwidth to reflect constraints that may in practice exist in the access network

5. Understanding applications

Bandwidth has no inherent value. Rather, it is an enabler of applications – everything from a simple text email to an on-demand 4K TV stream. Thus bandwidth needs can only be assessed by looking at applications – which are and will be used, how many people will use them, for how long, and how much bandwidth will be required?

In this chapter we consider both the current state of play, and the prospects for a wide variety of applications.

Wherever possible we have used third party sources for bandwidths and usage. In particular, we have used application providers' own estimates for bandwidth requirements where available. By way of summary, Figure 13 shows the Canadian regulator's assessment of required bandwidth for a range of applications. Note though that average speeds for a particular category of video may be exceeded for demanding content types such as live sports.[36]

Such actual figures are frequently lower than the headline figures sometimes cited when discussing bandwidth requirements.

Figure 13: Video & audio bandwidth per Canadian Radio-television and Telecommunications Commission[35]

Category	Application	Average speed (Mbps)	
		Down	Up
Video streaming	Netflix (Good)	0.63	0.02
	Netflix (Better)	2.22	0.07
	Netflix (Best)	5.94	0.18
	Rogers Anyplace TV	3.19	0.10
	CBC Television	4.42	0.11
	Global TV	0.97	0.03
	YouTube (HD 720p)	1.34	0.03
	YouTube (480p)	0.49	0.01
	YouTube (240p)	0.32	0.01
Audio streaming	Slacker	0.14	0.01
Live audio	Ottawa FM radio	0.03	0.00
	CBC radio	0.13	0.00
Real-time	Skype (w. video)	0.24	0.24
	Skype (audio only)	0.04	0.04
	Google Talk (w. video)	0.26	0.26

[35] CRTC, Communications Monitoring Report 2013, September 2013, except real-time applications which were not reported in 2013 and are from CRTC, Communications Monitoring Report 2012, 5 September 2012. CBC figures are the midpoint of the range in the CRTC report

[36] This is a result both of the rapid movement in the picture (which is inherently harder to compress than more static images) and because compression must be done in real-time. With more processing time available, a given video stream can be compressed more effectively than is possible in real-time.

Primary applications

Our focus for primary applications is video. The importance of video as a traffic driver is well known. Cisco's Virtual Networking Index estimates it is 74% of consumer traffic in Australia. (We note however that Cisco's traffic split seems to be inconsistent with some Australian ISPs' views that filesharing may represent one third of traffic – Cisco estimate this category at just 11%).

Figure 14: Australian 2013 consumer traffic by type, in petabytes (PB)[37]

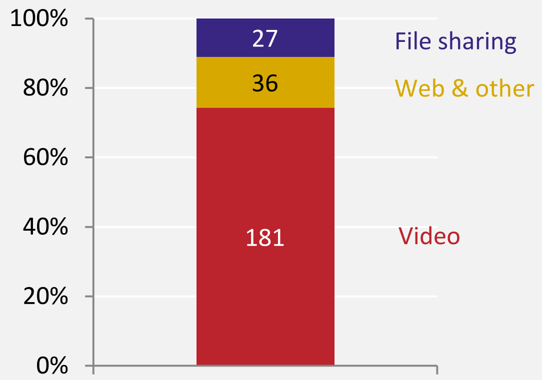

Within video the model considers: Internet TV (in various resolutions), YouTube and similar video; HD video calls; and streamed gaming.

Internet TV viewing

To forecast the bandwidth needs of primary Internet TV applications, we estimate:

1. The total amount of Internet TV, in minutes per user per month
2. How that viewing will be split between SD, HD and 4KTV

To estimate the total amount of Internet TV, we built up from current usage of such services in Australia.

Legal Internet TV viewing is still far from ubiquitous in Australia (and much less used than in comparator markets).[38] As of 2013, only 54% of those online had *ever* used an online catch up service, for example. However, 26% report watching TV content daily, and 59% say they would view more video if their internet connection speed was faster. [39]

One reason why Internet TV viewing may be relatively low is that internet connected TVs are still not widespread – they are in just 22% of Australian homes (up from 18% a year ago).[40] Of course, Internet TV can also be consumed on other devices such as tablets, which are in 37% of homes (up from 22%).

According to Nielsen, ABC is the broadcaster Internet TV site with the heaviest usage, with 1.3m users spending an average of 2 hours

[37] Cisco, *Visual Networking Index*, accessed April 2014
[38] See ACMA, *Online Video Content Services in Australia*, 8 October 2012
[39] Deloitte, *Vox Populi: State of the Media Democracy Survey 2nd Edition - Australia's media usage and preferences 2013*, 2013
[40] Oztam, Live *TV Viewing Drives Australian's Screen Use: Q3 2013 Multi-screen Report*, 7 February 2014

33 minutes in March 2014.[41] This corresponds to 11.7 minutes per internet user per month. (These figures are for access from a computer only, which represents approximately 60% of video consumption, with tablets and mobiles making up the other 40%).[42]

ABC is one of a number of other broadcaster Internet TV services - according to the ACMA, there were nine providers operating in Australia at the end of 2012, including services such as Foxtel and Fetch TV.[43] In addition, Vevo, a site providing professional music videos, is roughly as large as ABC, with an average of 10.2 minutes of viewing per internet user per month.[44]

We estimate total current Internet TV viewing (of the types tracked by Nielsen) at approximately 40 minutes via PC. Across all devices, this figure would rise to 65 minutes.

However, in addition to the video usage tracked by Nielsen, there are two other categories which it does not cover – video via VPN and adult content. Adult we treat below alongside YouTube, since we believe its video bit rates have more in keeping with that service than with professional Internet TV offers.

VPNs are used to access popular international Internet TV services such as Netflix which are not yet officially available in Australia. A VPN provides an intermediary US IP address to Netflix to bypass its geoblocking, and then passes the video traffic on to the consumer.

Reliable figures for Netflix usage are unavailable, though one estimate is that it has between 50,000 and 200,000 subscribers in Australia.[46] Certainly users' Google search volumes suggest that Netflix is getting considerable attention in Australia, despite not officially being available here – searches for the term 'Netflix' are 2.7x more frequent than those for 'Quickflix', for example (Figure 15).

Figure 15: Search volumes for internet TV terms[45]

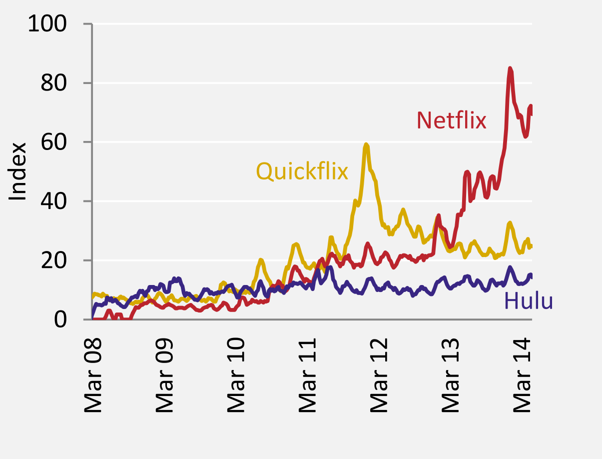

A typical international Netflix subscriber (a household) may consume 58 minutes of

[41] Nielsen, *Australian Online Landscape Review – March 2014*, 30 April 2014
[42] OZTAM / Nielsen, *Australian Multi-screen Report – Q4 2013*, 11 April 2014
[43] ACMA, *Report 1 – Online video content services in Australia: Latest developments in the supply and use of professionally produced online video services*, 2012
[44] Nielsen, *Australian Online Landscape Review – March 2014*, 30 April 2014
[45] *Google Trends* [accessed 8 May 2014]. Search volumes are for Australian users only
[46] Michael Body, "More Netflix pirates on board", *The Australian*, 3 March 2014

content per day from the service.[47] These figures would imply – very roughly – that Netflix via VPN adds a further 11 minutes of internet TV per online Australian. Usage of other international Internet TV services, such as the BBC's iPlayer, would be in addition.

Taking these various forms of Internet TV into account:

> We assume that in 2013, average Internet TV viewing is 80 minutes per user per month.

Clearly there is substantial potential for Australian Internet TV to grow from this low base. Increasing penetration of internet-connected TV sets, set-top boxes and tablets will all provide more opportunities to view. Further, those who already have access may increase their volume of viewing.

However, the potential for growth is not boundless. In the UK, the BBC's iPlayer service has been available since 2007, and the great majority of broadband connections can support streaming video. In this advantaged environment, iPlayer is now serving 198m TV programme requests per month.[48] However, while iPlayer is very well regarded, we estimate this viewing represents just 4% of the BBC's total, with the great majority still being via broadcast.

Thus our expectation is that most TV viewing (in Australia as elsewhere) will continue to be linear, though on-demand will grow appreciably.

> We assume that Internet TV viewing grows to be 25% of all Australian TV viewing by 2023 (equivalent to 47 minutes per day or 1,395 per month for all adults).

Note that there is the possibility that in future smarter DVRs (with ever larger hard-drives) will anticipate our content needs and record shows we are likely to want to watch even without being instructed in advance.[49] This will eliminate the need to deliver those shows on-demand via the internet. However we have not factored this into our forecasts.

SD, HD and 4K Internet TV viewing

A significant driver of Internet TV bandwidth requirements will be the resolution and format of the video. Main set viewing is

---

[47] BTIG, *Netflix Domestic Streaming Usage Up Low Double Digits Per Sub in January 2014*, 4 February 2014
[48] BBC, *iPlayer Monthly Performance Pack – February 2014*, 27 March 2014
[49] See for example: King's College London, *'Intelligent and green' iPlayer records favourite TV in advance and reduces internet traffic*, 17 June 2013

currently migrating from standard definition to HD, and may in time shift to 4K.

The proportion of households in Australia with a HD TV set grew to 74% in 2013 (with 48% having a HD TV service).[50]

We have assumed that the proportion of HD households is 74% in 2013, growing at 2% per year (excluding those who upgrade to 4K TV).

While demand for HD resolution is real, viewing of HD content is not that high. Not all programming is available in HD and even when it is, viewing is often still of the SD versions. In the UK FEH Media Insight analysed the viewing of HD channels (versus their SD equivalent) in HD households, and found that "the HD to SD viewing ratio has remained relatively flat at between 10% and 13% over the last 16 months".[51]

We have assumed that HD viewing represents 20% of Internet TV viewing in HD households in 2013, growing by 6 percentage points per annum, to 80% by 2023.

"4K" or "ultra HD" TV services have dramatically increased resolution. While a small number of 4K TV sets are now available, penetration is extremely low and anticipated to remain low in the short term constrained by factors such as the affordability and large size of sets and the lack of available 4K content. Futuresource forecast that penetration of 4K sets will reach 19% in 2020 in the US. [52] We have been unable to identify forecasts for Australia, and thus have based ours on Futuresource's US figures (with continuing growth to 2023).

Based on Futuresource's forecast, we have assumed that 4K penetration will grow from 0% in 2013 to 43% in 2023.

Those who have invested in a 4K TV may be particularly likely to seek out 4K content, but in practice there may be limited availability of such content.

We have assumed that viewing of 4K content will grow to 28% of all Internet TV viewing in 4K households by 2023.

[50] Ofcom *International Communications Market Report 2013*, 12 December 2013
[51] Farid Husseini, *High Definition Television: It's a must have, but is it also a must see?,* 11 July 2012
[52] Futuresource, *Ultra HD: The Content Perspective*, 27 June 2013

SD, HD and 4K Internet TV bandwidth requirements

The bandwidth required to stream SD video varies by service and device, although tends to be in the region of 1.5 – 2.5 Mbps. For example:

- The EBU estimate that the bitrate for an SD programme is 2.65 Mbps,[53] of which 1.8 Mbps relates to the video and 0.85 Mbps for the associated data
- According to Netflix, the average realised bandwidth of their streaming video services was between 1.0 Mbps and 2.9 Mbps in the US in March 2014,[54] and between 2.3 Mbps and 2.9 Mbps in the UK.[55] These averages will include some HD streams increasing the bitrate. However since it is a realised figure it also reflects factors which impair performance such as upstream congestion, home network congestion, distance from exchange etc. (decreasing the bitrate)
- ABC's iview has a requirement for a 1.1 Mbps or better internet connection[56]
- BBC's iPlayer uses 1.5 Mbps for SD streams, although lower bandwidths are used for some devices (for example, the BBC caps mobile phones at 800kbps over wifi networks)[57]

We have taken an approximate midpoint and assumed that SD video streaming requires a bandwidth of 2 Mbps in 2013.

(Note that the model works in 1 Mbps increments for downstream traffic, and thus all apps must be rounded to the nearest Mbps – in some cases this slightly understates true bandwidth, and in others overstates it. However, the net impact on total bandwidth requirements is minimal).

The bandwidth requirements for streaming of HD video content vary considerably, in part because of differences in format and resolution.

- Quickflix report that 4 Mbps is sufficient for HD[58]
- YouView (a UK streaming platform) recommends a minimum broadband speed of 3 Mbps for a service which can stream HD content[59]

---

53 Elena Puigrefagut for the EBU, *HDTV and Beyond*, 2012
54 Netflix, *US ISP Speed Index*, March 2014
55 Netflix, *UK ISP Speed Index*, March 2014
56 ABC, *ABC iview – catch up from ABC TV*, accessed March 2014
57 BBC, *iPlayer FAQs*, accessed August 2013
58 Quickflix, *Streaming of movies and TV enters the mainstream*, August 2013
59 YouView, *Broadband speed FAQs*, accessed August 2013

- BBC iPlayer requires 3 Mbps for HD content based on 720p25 (720p resolution as 25 Hz) [60]
- According to Netflix, 'Super HD' (in 1080p standard) requires between 5 Mbps and 7 Mbps[61]
- In 2012 the EBU estimated that the bitrate for an HD programme varies between 7.85 and 10.85 Mbps[62]

We have assumed the HD video requires 5 Mbps, an approximate mid-point between these estimates.

We assume this is a blended rate between 720p, 1080p and 1080i, the most common HD formats.

4K services have significantly larger bandwidth requirements. Again, estimates of these requirements vary:

- Netflix state that their "4K streams are encoded at 15.6Mbps"[63]
- Deloitte state that in 2014 4K TV will require 20 Mbps[64]
- *Satellite Executive Briefing* reports that the implementation of HEVC[65] will reduce requirements from 48 Mbps to "24 Mbps at the most"[66]
- Eutelsat suggest less than 26-40 Mbps for 4K at 50 fps[67] using MPEG-4, falling to 13 to 20 Mbps with HEVC[68]
- Docomo report that they have achieved 10 Mbps using HEVC[69]
- Encoding technology supplier ATEME state: "We're showing 4K at 60 Hz with excellent quality between 11 and 15 megabits per second"[70]
- According to video processer Elemental: "Today, 4K is simply bandwidth prohibitive for distribution networks at 18 Mbps to 20 Mbps required for H.264 compressed 4K content. HEVC shrinks bitrates required for 4K resolution potentially to under 10 Mbps"[71]

---

[60] BBC, *iPlayer FAQs*, accessed August 2013
[61] Netflix, *Netflix Super HD*, accessed September 2013
[62] Elena Puigrefagut for the EBU, *HDTV and Beyond*, 2012
[63] Netflix, *Letter to shareholders*, 22 January 2014
[64] Deloitte, *Survival of the fastest: TV's evolution in a connected world*, September 2013
[65] High efficiency video coding, a recent video compression standard now appearing in consumer equipment and codec software
[66] Elisabeth Tweedle, "4K TV: A technology push or a demand pull?", *Satellite Executive Briefing*, September 2013
[67] Frames per second
[68] Eutelsat, *Beyond HD, ULTRA HD*, 22 August 2013
[69] DigInfoTV, *Docomo demos HEVC (H.265) new video coding standard,* 19 February 2013
[70] Fred Dawson, "Underestimating 4K Potential Poses Mounting Risks for SPs", Screen Plays, 7 October 2013
[71] "HEVC goes beyond HD", TVBEurope, 4 June 2013

We have assumed that 4KTV requires 16 Mbps in 2013.

In our modelling approach we have not explicitly addressed 3D TV and have instead combined it with 4K TV. Note however the generally lower bandwidth requirement for 3D TV versus 4K. According to Netflix, 'best quality' 3D streams require up to 4.7GB per hour - this equates to a bandwidth requirement of 10.7 Mbps.

The upstream requirements for Internet TV are likely to remain minimal. CRTC found that even the highest quality Netflix video streams had an average upstream bandwidth of less than 200 kbps.[72]

We have assumed an average upstream bandwidth of 0.2 Mbps, for all modes of video (with no compression improvement).

Compression

The bandwidth for a given quality of video has fallen steadily over time, and will continue to do so, due to both improvements in compression techniques within standards and transitions to new standards (such as that from MPEG2 to H.264/MPEG4, and now on to HEVC). These in turn are enabled by ever greater processing power available to codecs – compression benefits from Moore's Law.[73]

Estimates for the rate of improvement in compression vary, although tend to be around 10%:[75]

- Sky in the UK suggest that the bitrate requirement for HD content has fallen from around 30 Mbps in 2002 to 10 Mbps in 2010[76], equating to an annual decline of around 13%.
- Zetacast suggest that video bandwidth requirements halve every seven years (see Figure 16), corresponding to an average annual decline of around 9%.

Figure 16: Required video bandwidth over time (adapted from ZetaCast)[74]

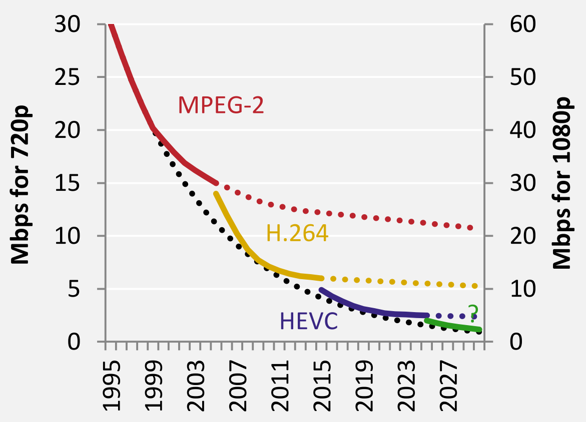

[72] CRTC, Communications Monitoring Report 2012, 5 September 2012
[73] A codec is the software which generates a video image from the digital stream of data carrying a compressed version of that image (or vice-versa)
[74] ZetaCast, Technical Evolution of the DTT Platform, 2012
[75] See also Brian Williamson, Anchor product regulation – retrospective and prospective, October 2013 for a useful discussion of changing bandwidth requirements for video and applications more generally
[76] Sky, Beyond HD Masters 2013, 2013

- According to Motorola, encoder technology evolution has resulted in a 10-20% annual improvement in efficiency[77].
- A 2006 University of Essex study predicted annual reduction in bitrates of 7%[78] (although with improving quality).
- Cisco, in their traffic forecasts, use a 7% annual reduction.[79]

> We have assumed an annual improvement in compression of 9% for SD, HD and 4K TV (and other forms of video).

YouTube and other non-TV video viewing

By some margin the leading player in this category is YouTube. It accounts for a large proportion of today's network traffic – according to Sandvine it represents 31% of peak downstream traffic in the Asia-Pacific region.[80] Since the usage and bandwidth characteristics of its traffic (and its peers') are materially different to that of Internet TV, we have considered it separately.[81]

According to Nielsen, the average online Australian consumed 9 hours and 35 minutes of video in Q4 2013, across all platforms.[82] As discussed above, 65 minutes of this comprises professional internet TV, leaving 8 hours, 30 minutes (or 510 minutes) per month for YouTube and other such services. Of this, we estimate YouTube comprises 92%, based again on Nielsen figures.[83]

In addition to this tracked usage, we need to account for adult video content. There is very little solid data on such usage, though we note that there are four adult sites in the Australian top 50.[84] SimilarWeb estimate that 7% of Australian web traffic is to adult sites.[85] We allow an additional 20 minutes for adult video.

> We have assumed the average adult watches 530 minutes of non-internet TV video per month from the home, growing to 1,000 minutes (33 minutes per day) by 2023

Required bandwidth for services such as YouTube varies significantly and is highly dependent on the resolution and format of the video stream:

---

77 Motorola, Opportunity and impact of video on LTE Networks, 13 May 2013
78 Ghanbari et el, Future performance of video codecs, 13 July 2006
79 Cisco, Cisco Visual Networking Index: Forecast and Methodology, 2012–2017, 29 May 2013
80 Sandvine, Global Internet Phenomena Snapshot: 2H 2013, 11 November 2013
81 We note that technically the two are similar, with both typically using Flash Video, albeit at different bit rates
82 OZTAM / Nielsen, Australian Multi-screen Report – Q4 2013, 11 April 2014
83 Nielsen, The Australian Online Landscape Review, February 2014
84 SimilarWeb, Website rankings in Australia [Accessed 9 May 2014]
85 Daniel Buchuk, "UK Online Porn Ban: Web Traffic Analysis of Britain's Porn Affair", *SimilarWeb Blog*, 27 July 2013

- As of 2010, a high resolution video of 1920x816 can require up to 3.4 Mbps, whereas a video of resolution 320 x 136 might need as little as 257 kbps[86]
- Ameigeiras et al's analysis of a sample of YouTube videos[87] found overall average bitrates of 510 kbps, and averages by resolution and format of 301 kbps (240p), 527 kbps (360p) and 840 kbps (480p). Over 99% of videos had encoding below 1.3 Mbps

We have assumed average downstream bandwidth of 1 Mbps.

Over time we would expect there to be both upwards and downwards pressures on the bandwidth requirement. We have assumed the upward trend in video resolution and associated bandwidth requirements will be offset by efficiency gains due to improvements in compression.

We have assumed an equivalent upstream requirement to internet TV of 0.2 Mbps (with no compression efficiency gains).

HD video calls

SD video calls (discussed below under secondary apps) are widespread. HD video calls are not, and are largely limited to the workplace. However there are signs that it will become a more mainstream activity - Skype now offers HD video calling on the iPhone 5 and fourth-generation iPad[88], for example.

We have assumed that the average user will make 30 minutes of HD video calls a month in 2017, growing at a rate of 20% per annum.

Like video more generally, the bandwidth requirement of video calls services varies depending on factors such as the resolution and format.

For one-on-one calls, a HD quality Skype video call requires 1.2 Mbps with 1.5 Mbps recommended.[89] Apple iChat requires up to 900 kbps for 'best' quality.[90]

Three, four and more way sessions may also increase the speed requirements further. For example, Apple state a 4-way video

[86] AdTerras Per Aspera blog, *Approximate YouTube bitrates*, 24 May 2010
[87] Ameigeiras et al, *Analysis and modeling of YouTube traffic*, 26 June 2012
[88] PC Mag, *Skype Update Adds HD Video Calling to iPhone 5, 4th-Gen iPad*, August 12 2013
[89] Skype, *How much bandwidth does Skype need?,* accessed September 2013
[90] Apple, *iChat system requirements*, accessed September 2013

conference at 'best' quality requires 1.8 Mbps for the initiator, and 300 kbps for 3 participants.[91]

Note that the lower rate of movement of 'talking head' video means that video calls tend to have much lower bandwidth requirements than the equivalent internet TV standard. The CRTC found that Skype and Google Talk video had an average speed of 200 – 300 kbps, considerably lower than SD video requirements of 1.5 – 2.5 Mbps.[92]

We have assumed an average initial requirement of 1.5 Mbps for HD video calls (upstream and downstream)

We would expect compression improvement similar to that of video more generally.

We note that WIK estimate that 25 Mbps is required to support HD video-communication,[93] although this appears extremely high and, given the much lower requirements of HD TV, presumably relates to a requirement of both very high video resolution and a large number of participants.

We also note that e-health may make use of HD videoconferencing. However, some medical video consultation works well on far lower bandwidth – a NIH study achieved good results using 3G macrocellular networks.[94] In Australia, Medicare Local recommends 1.7 Mbps for HD videoconferencing.[95]

HD interactive / Streamed gaming

'HD interactive' refers to activities where video is created 'on the fly' at a server and then streamed to the user, in response to user mouse-clicks and key strokes.

A prime example would be streamed gaming - Onlive and Gaikai (bought in 2012 by Sony for $380m) are examples of such services. From the perspective of games publishers, this approach reduces the risk of piracy, potentially cuts out intermediary retailers and eases patching.

Note that this is not the approach used by most online gaming today, where game state information is shared between players,

[91] Apple, ibid
[92] CRTC, *Communications Monitoring Report 2012*, 5 September 2012
[93] WiK, *Market potential for high-speed broadband connections in Germany in the year 2025*, 15 January 2013
[94] Locatis et al, *Video Medical Interpretation over 3G Cellular Networks: A Feasibility Study*, December 2011
[95] Medicare Local, *Is my Practice Technology Suitable for Telehealth?*, November 2012

but all video rendering is done locally. As we will see, that method requires very little bandwidth.

In practice there is limited usage of streamed gaming / HD interactive content today. However, if PC and console gaming moves towards a streaming model over time, it has the potential to become a more widespread form of video.

We have assumed that streaming gaming / HD interactive grows from nothing today, to 450 minutes per user per month by 2023.

In terms of bandwidth requirements, Microsoft recommended a minimum connection speed of 3 Mbps for Xbox Live gaming[96] (an early example of cloud gaming which has since closed). Even for the largest TVs, such services require 5 Mbps, though of course this is required as long as the game is being played, not simply while software is downloaded.[97]

In our model we have assumed the bandwidth requirement is equivalent to that of HD TV – 5 Mbps in 2013.

Upstream requirements are still relatively modest. We have assumed an upstream requirement of 0.2 Mbps.

Secondary applications

Secondary applications cover those activities which are amenable to multitasking – for instance, launching a movie download and then continuing with other activities, such as web use.

Like our treatment of primary applications, we have considered the characteristics of each application type, rather than the specific use case.

Cloud storage

Cloud storage refers to those applications which synchronise files stored locally with back-ups in the cloud. Such applications are already widely used today - Dropbox, for example, has 275m users globally, and 1bn files saved each day.[98]

Cloud storage services such as Dropbox tend to run in the background and use bandwidth as needed and when available. Dropbox downloads are performed at the fastest available

[96] Microsoft, *Xbox Live support*, accessed September 2013
[97] Onlive *Tech Specs*; Time, *Sony Bets On Cloud Gaming with Gaikai Purchase, but Don't Expect Drastic Changes*, 3 July 2012
[98] Dropbox, *Company Info* [Accessed 20 April 2014]

download speed, and uploads are automatically throttled to 75% of the maximum upload speed to prevent slowdown in browsing.[99] Furthermore, bandwidth requirements are eased by a lack of user awareness – if a file takes longer to upload, the user is unlikely to notice.

To estimate the time and bandwidth requirements of a cloud storage service like Dropbox, we have taken the average downstream and upstream throughput based on analysis by Drago et al[100] (1.26 Mbps and 0.54 Mbps respectively). We note that these are technically 'constrained' figures, in that they are a consequence of available bandwidth today. However, in practice they appear to meet users' needs – we are unaware of any complaints about the speed with which Dropbox uploads or downloads files. This is in contrast to loud complaints about the time to download new mobile OSs, for example.

We assume that throughput grows over time, driven by a 10% annual growth in users' speed expectations.

To calculate the implied minutes of use, we use Drago et al's estimate of total traffic per household of 0.34 GB per month (or 0.17 GB per individual).[101] Based on ACMA figures, we estimate penetration of cloud storage services such as Dropbox at 10% today,[102] and assume it will rise to 75% in 2023. We also assume an increase in the average number of shared devices per user from 1.2 today to 2.0 in 2023.

We note that some claim extremely high bandwidth requirements for cloud computing. WIK suggests speeds of up to 100 Mbps may be required.[103] However they assume files will be stored remotely but processed locally. To make this work high speeds are necessary – according to WIK "the typical LAN-feeling of about 100 Mbps would satisfy those needs".[104]

While such an approach is conceivable, we have not included it in the model. Firstly, we are not aware of any material usage of this technique today, even in markets with widespread superfast broadband. Secondly, this approach would be unusable if the user was offline, and problematic if they were on a wireless connection out of home. Thirdly, it would require high speed end-to-end, not

[99] Dropbox, *FAQs* [Accessed 20 April 2014]
[100] Drago et al, *Inside Dropbox: Understanding Personal Cloud Storage Services*, 14 November 2012
[101] Drago et al, *Inside Dropbox: Understanding Personal Cloud Storage Services*, 14 November 2012
[102] ACMA, *Cloud computing in Australia*, March 2014
[103] WiK, *Market potential for high-speed broadband connections in Germany in the year 2025*, 15 January 2013
[104] FTTH Council, *Q&A Webinar on WiK report*, 15 January 2013

just in the access link to the customer. Thus it would be vulnerable to network congestion, problems on the consumer's wifi and so on. Fourthly, it does not appear to bring material benefits over the approach taken by Dropbox, which is already being rapidly adopted.

We have above discussed cloud storage, but of course there are other kinds of cloud applications, in particular those where both storage and processing take place in the cloud. Examples include Gmail, Salesforce.com (for salesforce management), Freeagent (for accounting) and many others. Since these all use a web interface, our model captures them within its web usage assumptions.

Content downloads

As well as streaming video, consumers may also download content, such as TV programmes, films and music files. Video content will, in traffic and bandwidth terms, be the most important.

AHEDA report that there were 8m download purchases of TV and film (in addition 14m streamed purchases, which are included in our Internet TV category).[105]

If we assume these 8m purchases have an average length of 60 minutes of content - representing a mix of films and TV shows - this implies 2.6 minutes of content per month per online Australian. However, the great majority of those online do not consume content in this way – we estimate that in the region of 10% do so.[106] (By implication, those who do download content legally acquire 26 minutes of content per month).

We have assumed that the average content downloader downloads 26 minutes of content per month in 2013.

We have assumed that this will grow by 10% CAGR, although note that growth in the popularity of streaming video services such as Netflix, and improved macrocellular networks (4G), mean increasingly there are viable streaming alternatives to content downloads (both in and out of the home).

While today around 10% of internet users legally download content each month today, we assume this rises to 40% penetration by 2023. However, we note that international evidence suggests this may be a high figure for the long term. According to Ericsson research in several countries,[107] the portion of respondents who

[105] AHEDA, *Game of Thrones leads Oz (legal) downloads in 2013*, 17 April 2014
[106] Communications Chambers estimate based on proprietary ISP consumer research
[107] US, UK, Germany, Sweden, Spain, Taiwan, China, South Korea and Brazil

download movies and TV shows more than weekly has fallen from 38% in 2011 to 29% in 2013.[108] It seems likely that downloading is being replaced by streaming.

The bandwidth necessary to deliver this content is a complex question. Five minutes of SD TV content (compressed to 2 Mbps) would need five minutes to download with a capacity of 2 Mbps, or 1 minute to download at 10 Mbps – both are technically viable, the only issue is the user's urgency to receive the content. (A 1 Gbps connection could theoretically deliver the content in under a second, though in practice other links in the chain from the server to the user device might become constraints in this case).

We have assumed that the starting expectation is that content downloads at twice real-time (e.g. a 50 minute TV programme would take 25 minutes to downloaded), rising to four times.

Note that this is the minimum download speed, supported by the bandwidth allowed for content downloads by the model at periods of peak demand. However, any downloads occurring outside these peak periods in the month would be faster.

Peer-to-peer

P2P, and in particular BitTorrent, is a major driver of traffic, comprising both legitimate and illegal file transfer (though it is generally accepted that the latter represents the bulk of the traffic).

P2P downloading is particularly popular in Australia (Figure 17). We understand that P2P represents approximately one-third of downstream traffic in Australia. This compares to 5% in the US and 11% in Europe.[110]

IPAF[111] consumer research found that 11% of online Australians adults are "persistent illegal downloaders" and 14% are "casual illegal downloaders".[112] Given the challenges of asking survey respondents to accurately report illegal behaviour, it may

Figure 17: P2P downloads per capita (Index)[109]

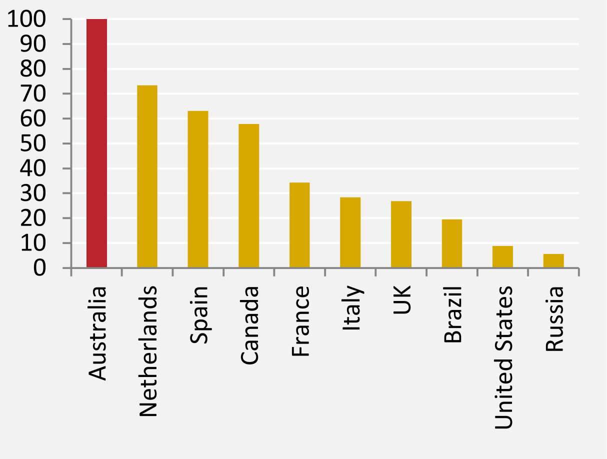

[108] Ericsson, TV and Media - Identifying the needs of tomorrow's video consumers, August 2013
[109] Irdeto, Piracy Report, 7 April 2014. Note that Irdeto's analysis focused primarily on US and UK made programmes which may have suppressed downloads in markets such as Russia. (It is not clear whether their sample included dubbed or only English-language versions)
[110] Sandvine, Global Internet Phenomena Snapshot: 2H 2013, 11 November 2013
[111] Intellectual Property Awareness Foundation
[112] Sycamore for IPAF, Adults and Online Movie and TV Piracy - Quantitative Wave 5 2013, 30 September 2013

be that these figures are underestimates.

Bittorrent is driving substantial usage in peak hours despite the implementation (since 2010) of a new protocol for BitTorrent file transfers, µTP. Designed to reduce congestion caused by BitTorrent in peak hours, µTP is a 'less than best efforts' protocol. It automatically reduces the bandwidth used for torrent transfers when it detects congestion – effectively, it time-shifts downloads to outside peak periods.

This means that BitTorrent is unlikely to cause quality of experience problems for other applications in use in a given household, but of course does not mean that other applications will not degrade the BitTorrent experience – if a household's access circuit is busy, torrents will take longer to download.

Thus as with other forms of downloading, torrents do not require a specific bandwidth – rather, the required bandwidth is a function of the expectations of the user. Obviously, if users expect to be able to torrent a movie in 5 minutes, far more bandwidth will be required than if they expect it to take five hours. As a practical matter, BitTorrent downloads can take hours rather than minutes. The About.com guide to BitTorrent says

> "Cable and DSL modem users can expect an average of 25 megabytes per hour [57 kbps], sometimes slower if the swarm is small ... On a good day with a big swarm, however, you can download … a 900MB movie within 60 minutes [2 Mbps].[113]

This is however an area where expectations are rising – some torrent sites now do offer streamed versions of files (that is, downloading in real time).

By their nature, P2P downloads are constrained by the availability and upload speeds of the relevant peers. Moreover, many ISPs throttle P2P traffic like BitTorrent (particularly in peak periods). If the constraint on transfer speed stems from peers' upload speeds, or is being imposed on such traffic at the core of the network, then increased bandwidth in the access network will bring no improvement in the user's quality of experience.

Assessing volumes of use of Bittorrent (and other P2P protocols) is challenging. As we have noted, consumer research in this area can be unreliable. We have used consumers' statements in the IPAF

---

[113] About.com, *How to download with BitTorrents*, August 2013

survey to estimate traffic, but found that our estimate fell far short of the one third of downstream traffic measured by ISPs (a presumably more reliable metric).[114]

We have therefore taken the approach of 'reverse engineering' consumer usage from total traffic. Based on total fixed broadband traffic per line and the one third ratio, we have calculated Bittorrent volumes, and from these estimated a number of minutes per downloader, on the basis of 30% of those online being downloaders. This 30% figure derives from 'rounding up' the proportion of respondents reporting using P2P in the IPAF research.

This approach results in approximately 1,200 minutes of video content downloaded per month per Bittorrent user (or 40 minutes per day). This level of consumption is somewhat greater that seen for Netflix users (which consume 58 minutes per *household*). It seems plausible that Bittorrent users may download more content than they in fact consume – extended download times and zero marginal cost likely encourage users to request (say) multiple episodes of a series, even if they are unsure they will subsequently watch them. By contrast, for a streamed service such as Netflix, minutes of content downloaded and consumed are likely to be similar.

In the short term we assume continuing strong growth in Bittorrent, at 25% per year. However, we further assume that this growth slows, with volumes reaching saturation in 2023, by which time just over 100 minutes of content per user per day will be downloaded.

In summary:

> We have assumed that 30% of individuals use P2P, rising to 40%, with average downloads of 1,000 minutes of content per month (40 minutes per day) in 2013, rising to 3,166 minutes in 2023.

Today, it is the SD versions of such P2P files that are more popular, with (according to one study) 92% of users choosing them.[115] Note that this preference is not necessarily as a result of the constraint of the individual user's bandwidth. Rather, it is the nature of P2P networks that if many users choose one format, that format will become more attractive to other users, since there will be more uploaders available to those other users. We have assumed that over time users will shift to higher resolution versions of content.

---

[114] Note that there was material variance in the ratios reported between ISPs – the reasons for this are not clear
[115] Petrus Potgeiter, *High-definition content and file sharing networks*, 18 November 2012

In determining bandwidth required, we have assumed that:

Initially users will (on average) expect P2P download in twice real time (120 minutes for a 60 minute programme), improving to in real time in 2023.

Note that this average of 'twice real time' reflects a blend of occasions when the user wants the content in real time, and other occasions when, say, the user is downloading an entire series, to watch over time later.

Non-HD video calls

As discussed above, we have assumed that HD video calls are a primary application, unlikely to be multitasked with Internet TV viewing or video gaming. However, we treat SD video calls as a secondary app, allowing the multi-tasking of such calls with these activities.

Nielsen report 25% of those online use Skype each month in Australia, with PC usage of 101 minutes per user.[116] Grossing up to include usage via mobiles suggests a total of 134 minutes per user. Allowing for other VoIP services (and averaging across all online adults) we have therefore assumed:

The average individual has 40 minutes per month of non-HD video calls in 2013, growing to 400 minutes by 2023.

We believe the assumption of 400 minutes in 2023 is conservative, even allowing for the possibility of video windows being left open for telepresence. This monthly usage is equivalent to combined fixed *and* mobile consumer voice minutes per capita today.[117]

For one-on-one video calls, a high quality Skype video call recommends 500 kbps[118] downstream and upstream capacity. As with other forms of video we would anticipate 9% annual compression improvements. However, as a practical matter we set bandwidth for such calls at 1 Mbps and hold it constant.

Software downloads

Software downloads are particularly challenging to address in the context of required bandwidth. The files involved can be large –the

---

[116] Nielsen, *The Australian Online Landscape Review, February 2014*

[117] Total outbound voice traffic is 215 minutes per person per month, or 430 minutes both-way traffic (ie including the call recipient). However, this figure includes business traffic. Ofcom, *International Communications Market Report 2013*, 12 December 2013

[118] Skype, *How much bandwidth does Skype need?,* accessed September 2013

average size across the launch range of games for the Xbox One[119] and PS4[120] platforms was 17.2GB (with NBA 2K14 on the PS4 a huge 41.8GB). Apple's latest OS X Mavericks, while smaller, is still 5.3 GB.

We will use Apple OSs and console games as our case studies, since we believe they represent significant drivers of the bandwidth requirement stemming from software downloads. In addition to new purchases, there may also be a steady stream of smaller patch files (though as we will see, these are far less demanding from a bandwidth perspective).

OS releases can lead to clear spikes in total network traffic, and downloading can be impractical for some consumers - indeed, when it released OS X Lion, Apple offered customers the option of bringing their computers to an Apple Store to transfer the necessary files.[121]

However, downloads are relatively infrequent. Telesyte report annual digital sales for full games (ie excluding mobile) of $265m for 2013.[122] On the basis of a unit price of $60, this suggests 4.4m downloads, or just over 0.5 per household per year.

One reason for infrequent games downloads is that there is still a strong preference for physical copies – while 63% of households have console games, 73% of adults report preferring physical over digital copies, and this implies only 17% of households download console games today.[123] Slow downloads are a factor in this preference, though earlier release dates for physical copies, desire for the manual and a physical copy for backup and the ability to trade-in are all factors too.[124]

Turning to OS updates, Apple updates its computer OS roughly annually, and has approximately a 23% share in Australia, implying 0.23 downloads per year per household (assuming 1 Mac per home).[125] For whatever reason, new Windows OSs appear to be much less likely to be downloaded.

As with other forms of download, determining the speed required is challenging. Consumers currently experience long download times – for console games overnight downloads are not uncommon. For

---

[119] Game Rant, *Xbox One Launch Lineup Install Sizes Revealed*, 11 November 2013
[120] Eurogamer, *All the PS4 launch games and their install sizes*, 13 November 2013
[121] Gigaom, *Lion download too fat? There's Apple store (WiFi) for that*, 21 June 2011
[122] VG24/7, *Retail strong and digital growing, Sony and Microsoft both pleased – NPD Australia*, 12 February 2014
[123] IGEA & Bond University, *Digital Australia 14*, 21 October 2013
[124] Deloitte, *Media Consumer Survey 2013*, 30 April 2013; Hybris, *PC and Console Gaming: The end of the DVD?*, 9 October 2012
[125] *Statcounter*

Mountain Lion, some users reported downloads in the range of 1 to 3 hours,[126] others as quick as 30 minutes.[127]

In an ideal world, such downloads would of course be instant. However, even a 1 Gbps connection would need over a minute to download a 9 GB console game.

Software developers are well aware of these challenges, and are developing workarounds. For instance, Microsoft is enabling the pre-loading of games prior to their release date, with authorisation to play issued on the day of release. [128] The new Xbox allows gamers to start playing a new game while the download is still in process. According to Microsoft "Once the required data – a fraction of the entire game – is on [a customer's] hard drive, they can jump into the action while the rest of the game finishes downloading in the background".[129] The recently released PlayStation 4 has the same capability.[130] The PS4 also wakes from standby to begin a download if a purchase for it is made from another location – for instance through a mobile device.[131] Thus the game can be downloaded before the purchaser returns to the console.

Furthermore updates may take place outside of peak hours, as they already do with Windows. Microsoft's Windows Update client "downloads updates using idle bandwidth. This technology ensures that Windows Update downloads only when no other active download is in progress on the computer. This allows you to smoothly carry on day-to-day activities even while updates are being downloaded in the background."[132]

A further transition that may impact the bandwidth requirement for game downloads is a potential shift to streamed gaming, discussed above. This may substitute for a certain number of downloads, though we have not factored this into our forecast.

We assume:

- The average console game size is 7GB for generation 7 devices (PS3, Xbox) and 17 GB for generation 8 devices (PS4, Xbox One)

---

[126] Tuaw, *How long will it take to install Mountain Lion?*, 25 July 2012
[127] Mashable, *Apple OS X Mountain Lion: The Good, the Awesome, the Could-Do-Better*, 25 July 2012
[128] IGN, *Ask Microsoft anything about Xbox One*, 5 August 2013
[129] Engadget, *Xbox One has play as you download functionality similar to PlayStation 4*, 16 July 2013
[130] Engadget, *PlayStation 4's UI and inner workings detailed: No more booting, games download as you play them*, 20 February 2013
[131] VentureBeat, *How PlayStation 4's play-as-you-download feature works*, 12 November 2013
[132] Microsoft, *Windows Update Explained*, September 2008

- Five downloads per year per using household, growing from 7GB to 17 GB in 2020 when generation 8 consoles are at maturity
- A 12% CAGR in filesize from 2020, based on an 8 year console generation cycle
- Initially only the 20% of households that are 'high use' download such files. Over time, a further 20% of households become downloaders
- By 2023, 90% of the download is pre- or post-loaded (growing from 0% in 2013)
- An arbitrary assumption that download expectation is 60 minutes, falling to 20 minutes in 2023

Note that there is one important conservative aspect of our approach to software downloads. The model calculates the average speed required to deliver a file of a given size in a given time, and assumes that this bandwidth must be available *throughout the download.* However, even if that bandwidth was temporarily unavailable during the download (for instance, due to some brief page loading for simultaneous web surfing), this needn't mean the file failed to arrive in the given time, if excess bandwidth was available at other times during the download – in effect, the download could 'catch up' during these periods of excess bandwidth. Thus our model results somewhat overstate the bandwidth impact of software downloads.

Mobile OS Downloads

While the need is only occasional, operating system (OS) updates for mobile devices such as smartphones and tablets can be sizeable and obligatory. They are also widespread.

Apple's most recent iOS 7 for iPhones and iPads was 752MB.[133] (Google's Android Jelly Bean update for the Samsung S4 smartphone was considerably smaller at 151MB).[134]

We have assumed a mobile OS update download size of 700MB, growing at 5% each year. We further assume that such updates are annual.

[133] The Telegraph, *iOS 7: how to install the new software for iPhone and iPad*, 19 September 2013
[134] Technostop, *Samsung galaxy S4 Jelly Bean 4.3 Update*, 3 August 2013

OS updates will apply to all those with smartphones (around 64%[135] today, which we assume grows to near ubiquity) and tablets (we assume the average user has 1.4 of such devices to account for tablet ownership, growing to 1.8).

Expectations of the download time are again a key driver for bandwidth requirements.

We have assumed a 30 minute download time in 2013, falling to 15 minutes in 2023

We note that current download times may be much longer than 30 minutes, with the recent iOS download taking "many hours" for some.[136] Particularly immediately after release, server capacity (not access network bandwidth) appears to be the key constraint for iOS downloads.

Content Upload

By some margin, the most popular site for photo sharing is Facebook.[137] Facebook users upload 350m photos each day. However, this works out to just 10 photos per month per user (given that Facebook has 1.06 billion active monthly users).[138] The average size of each photo is just 107 KB, likely because the Facebook app compresses photos before uploading them.[139] This implies a total upload per user of just over 1 MB per month – even at 1 Mbps, this would upload would take just 8 seconds. Note however that sites such as Flickr, though far less widely used, likely receive larger photo files.

Video drives heavier uploads, although this is not a widespread activity. According to 2011 research, just 6% of online Australians had uploaded a video in the past month.[140] More recent research in other markets tells a similar story – as of 2013, only 27% of US[141] and 22% of UK[142] internet users had *ever* uploaded video.

YouTube's reach in Australia is 55%, though of course this includes individuals who only view.[143] Globally, YouTube receives 100 hours of video every minute and has over 1 billion visitors per month,

---

[135] Kosmosion, *Smartphone Usage In Australia*, 2013
[136] The Guardian, *iOS 7 download delays frustrate Apple customers*, 19 September 2013
[137] Mary Meeker, *Internet Trends*, 29 May 2013
[138] Facebook, *Form 10-K*, 1 February 2013
[139] Communications Chambers analysis of figures from Dhruba Borthakur, *Petabyte Scale Data at Facebook*, November 2012
[140] ACMA, *Digital Australians - Expectations about media content in a converging media environment,* October 2011
[141] Pew Research Centre, *Online Video 2013*, 10 October 2013
[142] Ofcom, *Adults media use and attitudes report 2013*, April 2013
[143] Nielsen, *The Australian online landscape review – February 2014*, 26 March 2014

implying that the average visitor uploads 0.25 minutes of video per month. (In practice, many will upload none at all in any given month).[144]

For modelling purposes, we assume that initially 20% of internet users are monthly uploaders, at an average of 5 minutes per month. Over time, the number of users increases to 40% and average minutes to 15.

This volume of video uploads today is well beyond current YouTube volumes - we have deliberately overestimated to allow for other forms of uploads (such as email attachments) for which we have not been able to identify usage data to underpin specific assumptions.

We have also not been able to identify data on the average bitrate of uploaded video. YouTube transcodes received video to a variety of formats to enable viewing at lower bandwidths, but the original file may have been very high resolution. The iPhone 5 records video at 17 Mbps.[145] However, by default iPhones compress video substantially before uploading to YouTube – 1 Mbps is a typical bitrate.

As with all upload and download applications, the choice of ‘required’ bandwidth is arbitrary. Note however that the upload itself is not the time-consuming element of making a video available on YouTube. YouTube’s processing of a video can take anywhere from a few minutes to several hours, particularly for higher resolutions.[146]

- We have assumed that uploaded video runs at 2 Mbps, and that it should upload 1x real time (i.e. 1 minute to upload a 1 minute video).
- Over time, we assume volumes of video, bitrates and upload expectations all increase.

Web surfing

Browsing or interacting with web pages is, and is likely to remain, much the most important internet activity by time spent.

According to Nielsen, the average internet user in Australia spent just under 37 hours (or 2,211 minutes) online via desktops and

[144] YouTube, *Statistics* [Accessed 20 April 2014]
[145] GSMArena, *Apple iPhone 5 vs. Samsung Galaxy S III: All rise*, 25 October 2012
[146] YouTube, *Video stuck in processing*

laptops in February 2014.[147] However, this figure includes video usage which we treat elsewhere. Excluding this leaves 1,868 minutes of web usage via desktop and laptop. Today 29% of Australian page views are via mobile devices,[148] and incorporating this usage brings web time to 2,615 minutes, or just under 1½ hours per day. (This will include out-of-home usage however).

Data from Nielsen suggests the average time online via desktop and laptop is reasonably stable, generally between 36 and 40 hours per user per month with no clear growth.[149] ComScore data on the time spent browsing on a computer[150] actually suggests a year-on-year decline (of up to 20%), although this is likely to be due to a shift away from PCs and towards mobile platforms. Allowing for this additional usage we estimate historic growth in total web time per person at 5%. Going forward, to allow for this figure reaching saturation, we have assumed a CAGR of 3% between 2013 and 2023.

The bandwidth required for surfing is a complex question. In practice, the demand is 'spiky' – potentially no traffic flow at all while a page is being read, with a sudden spike in need when the next page is accessed.

This means an individual using the web will require more capacity than the traffic consumed might suggest.

Figure 18: Web usage and page loads (illustrative)

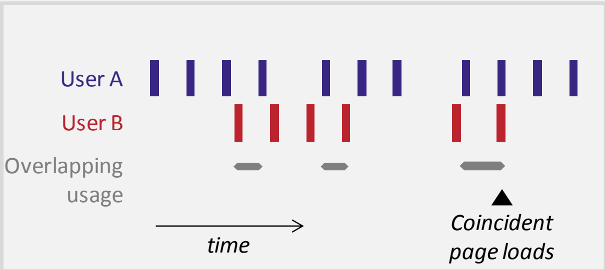

For our purposes it is the peaks that matter. If adequate bandwidth is not available to load that next page in a reasonable time, the user will be frustrated.

In practice, the experience of web-surfing is driven at least as much by latency (the round-trip time for a request for data and the response) as it is by bandwidth (how much data can pass simultaneously through a particular element of the network).[151] Indeed, beyond a certain point increased bandwidth makes relatively little difference. For instance, according to one study, increasing bandwidth from 1.5 Mbps to 5 Mbps might

[147] Nielsen, *The Australian Online Landscape Review*, February 2014
[148] Statcounter [Accessed 20 April 2014]
[149] Nielsen, *The Australian Online Landscape Review*, various reports from May 2013 to February 2014
[150] Referenced in Ofcom, *International Communications Market Report 2013*, 12 December 2013
[151] Note that different access technologies have different latencies – generally mobile will have the highest and FTTH the lowest

only reduce page load times by only 12%.[152] Others have found diminishing returns from 2 Mbps onwards.[153]

A feature of many pages is that they will load some content immediately ('above the fold' content which is visible when landing on the page), with other content (which sits below the fold) delivered later. This helps to smooth the 'spikiness' of the bandwidth requirement, improving the user experience.

We have taken 2 MB as the maximum required pre-caching page weight for above-the-fold content, and assumed bandwidth must be provided to download this in 3 seconds

Our 3 second assumption compares to a current Australian average of 3.9 seconds.[154]

We have also assumed:

- Total page weight rises to from 1.4 MB[155] in 2012 to 5.9 MB in 2023 (which increases how long the provided bandwidth must be sustained)
- Average time spent on a page is 77 seconds[156], growing at 2% per year (to reflect pages offering more 'below the fold' content)
- The downstream to upstream ratio for web traffic is 9.28[157]

Note that the above assumptions (in particular page weights) are based on the web as accessed from fixed devices. However, as discussed above, 29% of page views are already from mobile devices, and rising fast. In some cases such devices will receive the standard (heavy) web page, but increasingly they receive a light version designed by the site owner for the limits of mobile screen size and data caps. US Cellular, for instance, estimates that that the typical web page viewed on a smartphone has a weight of 180KB, roughly one-eighth of our overall assumption.[158] Such lower mobile weights will greatly reduce the average page weights for surfing, but this effect is not factored into our forecast.

---

152 NCC Group, *Will faster user bandwidth fix your website performance woes?* 5 June 2013
153 Ilya Grigorik, *Latency: The New Web Performance Bottleneck*, 19 July 2012
154 Akamai, *State of the Internet, Q4 2013*, 23 April 2014
155 Based on the weight average page weight of top 10 Australian websites (YouTube excluded, since it is treated separately in the model)
156 Communications Chambers analysis of Nielsen data, February 2014
157 Sandvine, *Global Internet Phenomena Snapshot: 2H 2013*, 11 November 2013
158 US Cellular, *Monthly Data Usage Estimate* [accessed 20 May 2014]

Low bandwidth applications

'Low bandwidth' is a catch-all bandwidth reservation for apps such as radio streaming and e-metering that are unlikely to make a material difference to household bandwidth requirements.

We have assumed that, in aggregate, low bandwidth applications require downstream bandwidth of 1 Mbps, and 1% occupancy during busy hours (growing to 3% by 2023).

We have assumed an upstream bandwidth of 0.2 Mbps, growing to 0.4 Mbps by 2023, and with occupancy rising from 1% to 50%. (We have assumed higher upstream traffic to allow for the constant flow of e-metering).

This category includes the following applications – note that these are unlikely to all be used simultaneously, meaning that their required bandwidth is likely to be well within the 1 Mbps downstream 'envelope':

Audio streaming

While we have not explicitly addressed audio streaming in our model, the required downstream bandwidth is low and likely to remain low. Most 'CD quality' audio requires around 128 kbps, and even extremely high quality streaming audio rarely requires more than 320 kbps.

Online gaming (excluding streamed gaming)

Online (non-streamed) gaming generally requires a surprisingly small amount of downstream capacity. For example, one of the most popular online games, *StarCraft II*, requires bandwidth of just 2-3 kbps[159], *World of Warcraft* 2 kbps and even high intensity *Halo 3* requires 60 kbps.[160]

Voice calls

These require 100 kbps or less.[161]

Smart metering

Bandwidths are often in the tens of Kbps or less. Italy's 30m smart meters use a bandwidth of 2.4 kbps.[162]

---

[159] Jose Saldana and Mirko Suznjevic, *Online Games: Traffic Characterization and Network Support*, 10 January 2014
[160] Andreas Petlund et al, *Network Traffic from Anarchy Online: Analysis, Statistics and Applications*, 22 February 2012
[161] Skype, *How much bandwidth does Skype need?,* accessed September 2013
[162] Sergio Rogai, *ENEL Telegestore Project*, presentation to Steering Committee of the Energy Efficiency 21 Project Ad Hoc Group of Experts on Energy Efficiency Investments for Climate Change Mitigation, 31 May 2006

Body telemetry

E-health may require monitoring of an individual's vital signs. However, while this is high value, it need not be high bandwidth. For instance, paediatric ECGs need "bandwidths less than a standard phone call".[163]

Remote working via VPNs and desktop sharing

Citrix estimate that their popular XenDesktop virtual desktop / remote working service requires an average bitrate of just 78 kbps.[164] Citrix report that PC-based desktop sharing (via WebEx) needs 43 Kbps on average, with a peak need of 598 Kbps.[165] Note that remote working may well involve other types of use, such as video calls, extensive web use and so on, but this usage is reflected within those applications.

Security cams

While still relatively rare in the home environment, companies such as Dropcam now offer internet-based video monitoring cameras. These typically use 0.2 Mbps (upstream) and occasionally 0.5 Mbps.[166] However, this bandwidth is generally only required when the camera is accessed remotely from a browser. It seems less likely that such cameras will be in use when occupants are themselves in the home making use of the internet, and as such these cameras are unlikely to contribute to peak bandwidth requirements.

Summary

As set out above, we have made detailed assumptions about volume of usage and bandwidth requirements for a wide range of applications over time. The results for 2023 are summarised in Figure 19 – note the log scale on the horizontal access. The highest bandwidth application, software downloads at 32 Mbps, has relatively brief demand, at 20 minutes per month. 4K TV, at 6 Mbps, has more sustained demand, at 165 minutes per month.

Figure 19: 2023 usage and bandwidth by app

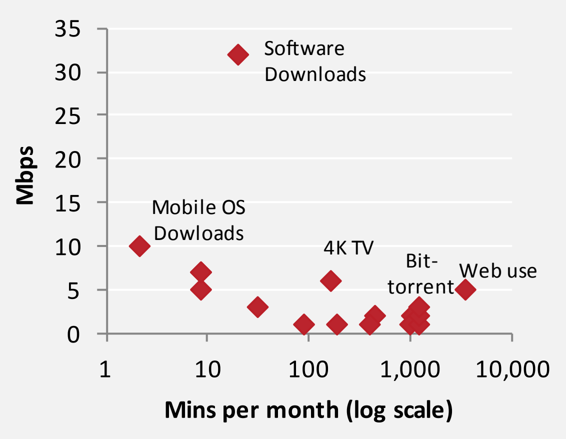

The heaviest application - as measured by time - is web usage, with 3,514 minutes per

[163] Harrop and Armitage, *Quantifying the Broadband Access Bandwidth Demands of Typical Home Users*, December 2006
[164] Citrix, *XenDesktop Planning Guide: User Bandwidth Requirements*, 28 July 2010
[165] Cisco, *Cisco WebEx Network Bandwidth*, June 2013
[166] Dropcam, *How much bandwidth does Dropcam use?*, 11 May 2012

month per adult (or 1.95 hours per day). While not the highest, its bandwidth requirement is still appreciable at 5 Mbps.

The applications that are most significant for peak demand are, broadly, those to the right or to the top of Figure 19. Heavily used apps such as surfing and internet TV create a 'base load' of bandwidth requirement. Demanding, briefer applications (such as mobile OSs) then have a significant chance of adding on top of this base load to create high app stacks that set peak bandwidth demand for brief periods in the month.

As Figure 19 shows, it is a relatively small set of applications which have high bandwidth requirements. This is because most applications are constrained by the recipient individual's ability to absorb information.

Broadly, we can categorise applications based on their urgency, into 'faster than real time', 'real time' and 'slower than real time'. A typical real time application might be internet TV, a radio stream or a cloud-based service, where the data is being consumed as it is downloaded – a consumer is actively paying attention to the stream in question. In such a case, the bandwidth requirement is set by the medium in question. For example 4K requires 6 Mbps (in 2023). In effect this sets an upper bound for any real time application (since it is the richest media type).

Figure 20: Urgency and richness of applications

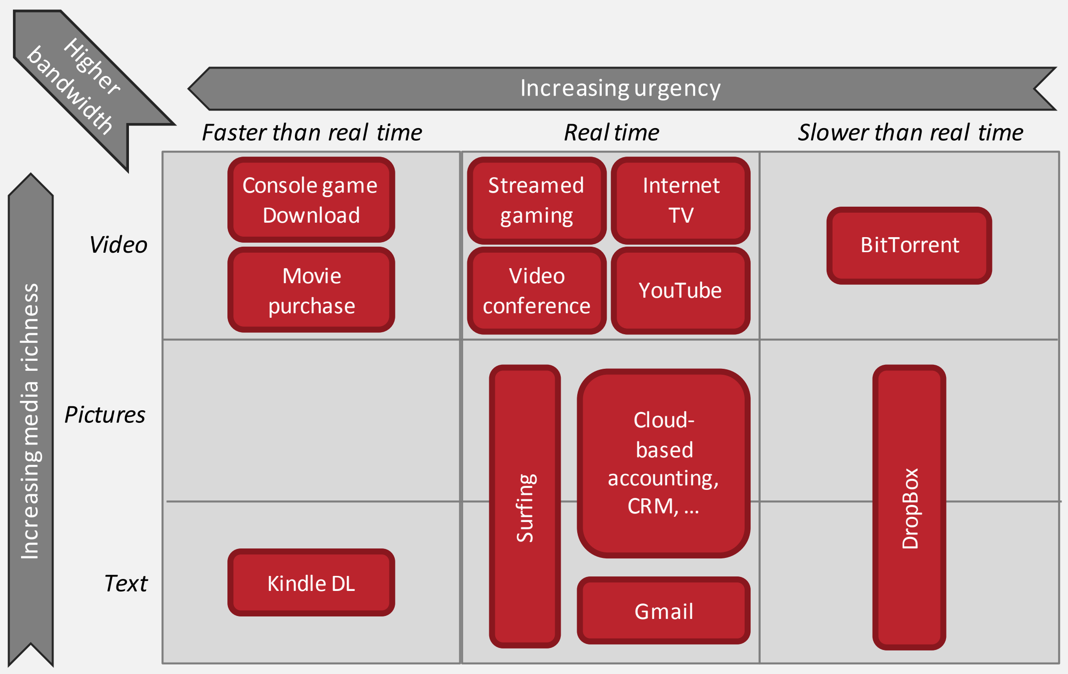

Examples of applications that may be ‘slower than real time’ are Bittorrent and Dropbox, where the download may take longer than the consumption.

There is then a limited set of applications where there is a need for the download to be quicker than the consumption – that is, cases where the media is pre-loaded for consumption later. Depending on the media in question, this need not require high bandwidth. For instance, a Kindle book is pre-loaded, but because the media in question (text) is so simple, file sizes and hence required bandwidth are low.

However, for rich media files such as purchased video or console games which are preloaded, required bandwidth can be significantly higher than that for any real-time application.

We now turn our attention to bandwidth demand results implied by our application assumptions.

6. Results

In setting out the results of the model, we first describe the metrics we will use to describe bandwidth demand. Secondly, we set out the base case results of our model. Thirdly, we investigate whether demand is driven by applications with positive externalities. Fourthly, we provide sensitivities to our base case results, to illuminate the wider range of possible outcomes, and to highlight the most sensitive input assumptions. Finally, we explore the basis for the differences between our conclusions and those of other forecasters.

While the model provides figures over time, for simplicity we will generally focus on 2023.

In interpreting these results, it is important to note what they are and are not. They are a forecast of technical bandwidth demand – that is, the actual bandwidth used by a household. This is not necessarily the same as the bandwidth that a household might be willing to pay for, which could be more or less. In particular we note that headline bandwidth speeds are an important competitive marketing tool for broadband providers, regardless of whether a given household is able to make regular use of the full speed offered.

Describing bandwidth demand

In considering household bandwidth requirements, we were struck that there is in fact little vocabulary to meaningfully describe such requirements. Commentators frequently make statements of the type that 'a typical household will need X Mbps'. However, such statements are incomplete in two ways.

Firstly, there is enormous variation in household needs, so the requirements of a typical household may not be particularly illuminating.

Secondly, the duration of requirement is crucial. If a consumer's peak need in a month is (say) 20 Mbps, but that demand lasts for just one minute per month, that is a very different situation from having a sustained demand for an hour per day of 20 Mbps.

In the former case, the consumer may be unwilling to pay for an upgrade to receive 20 Mbps – they would likely rather just wait the minute to begin one of the activities, or perhaps live with degraded video quality[167] or slower page loads in that minute (assuming their current bandwidth is sufficient for the rest of the month).

In the latter case, faced with an hour per day of degraded experience, the consumer might be far more likely to upgrade to 20 Mbps.

Thus we set out the results of the model in terms of 'X-minutes-excluded-monthly demand'. A demand of 20 Mbps on a 1 minute excluded basis means that if we ignore the household's single highest bandwidth minute of usage per month, 20 Mbps would be enough to fully satisfy the household's bandwidth needs for the rest of the month. (The ignored minute could require 25 Mbps or 100 Mbps, but in either event this is set aside).

In presenting our results, we primarily focus on '4 minutes excluded monthly' bandwidth.[168] This choice is arbitrary. Tightening to '1 minute excluded monthly' increases the required bandwidth, loosening to 10 minutes would decrease it.

Some commentators believe that the correct metric is zero minutes - that is, no matter how brief the demand, the network should be built to accommodate it.[169] However, infrastructure is almost never built to accommodate extreme cases – roads are not built to run without congestion at the peak of rush hour, reservoirs are not built to provide the full demand for water in a drought, and telecoms networks are not built to meet all demand at all times. The reason is simple – the cost of doing so would greatly outweigh the benefits.

For telecoms, 'five nines' is often said to be 'carrier grade' reliability. Five nines mean that availability is 99.999% (or all but 0.4 minutes per month). The significance of 'carrier grade' is that carriers are presumed to need far higher reliability than their customers. An outage for a carrier can affect many customers at once, and is obviously more serious than an outage at a customer premise affecting only that one customer.

We also note that degraded internet performance for an end user can be caused by many things – poor performance by the client device, a congested wifi or macrocellular network, congested

[167] Most video over the internet uses adaptive bitrate streaming. With ABS, the streamed picture quality adjusts automatically in light of available bandwidth. Thus, in the face of congestion, the viewer might receive SD rather than HD video for a period, but the stream would continue without interruption
[168] Roughly equivalent to excluding one peak minute per week
[169] See for instance Dr. Alessandro Monti [WIK Consult], *Market potential for high-speed broadband connections in Germany in the year 2025*, 15 January 2013, which says "Optimal user experience has to be guaranteed at any time. No limitations on usability or function"

transit links, problems at the application server and so on. This suggests that the benefits of extremely high performance in the access link may be 'lost in the noise'.

Against this background, we have taken the view that '4 minutes excluded monthly' is a reasonable threshold for a household's bandwidth – it is equivalent to 99.95% full availability, and as we have noted, even within the four excluded minutes the most likely adverse consequence is a moderate degradation of video resolution rather than anything more serious. By contrast, 'five nines' availability may mean a complete loss of service for 0.001% of the time, not a degradation.

Model results

Usage profiles of individual households

As we have described, the model works bottom up, building from an individual's usage of a range of applications, combining these usages to get that individual's overall usage profile, and in turn combining these to create the usage profile of a particular household type. These household profiles are then combined according to nationwide household mix to get to a national profile of usage.

In all, the model develops usage profiles for 192 different household types:

Figure 21: Household types

16 Demographic types	X	*4 Usage intensities*	X	*3 Main TV types*	=	*192 HH types*
1 adult, 2 adults, 1 adult, 1 child …		High/Med with software DL & video UL/Med/Low		SD/HD/4K		

Figure 22 shows the downstream usage profile for three sample household types. Minutes are ranked from those needing the highest bandwidth to those needing the least. (The total number of minutes is 7,200, representing four busy hours per day over 30 days).

Figure 22: 2023 usage profiles, three sample household types

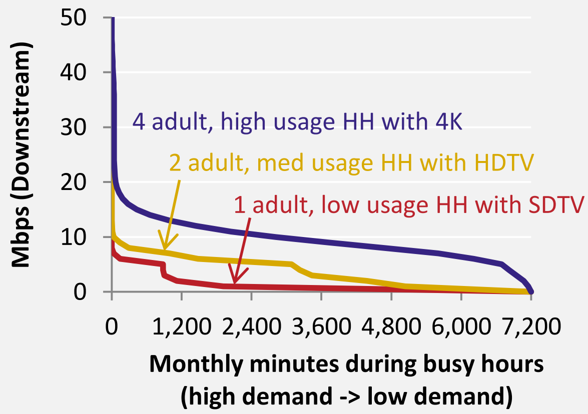

In the 1 adult, low usage household, the broadband connection is idle for much of the month, requiring 0 Mbps. At the other extreme, the 4 adult, high usage household with a 4K TV has demand almost constantly during the busy hours, frequently above 10 Mbps and occasionally above 20 Mbps.

Note the 'shelves' on the graph – these represent the required bandwidth for a particular common combination of apps (or one heavily used app in isolation). For instance, the 1 adult household has a shelf at 5 Mbps – this represents the bandwidth required for substantial web surfing (by itself) – a common 'app stack' for a single person household.

By zooming in on the busiest minutes, we can assess the '4 minutes excluded' demand for each of these sample household types (Figure 23). We read off the bandwidth that is sufficient to serve the peak demand of the fifth most busy minute (the '241st second') and all less busy periods. In the case of the 4 adult household, this is a required bandwidth of 46 Mbps, though note that in the excluded peak minutes, the bandwidth requirement reaches over 60 Mbps.

Figure 23: 2023 usage profiles, three sample household types

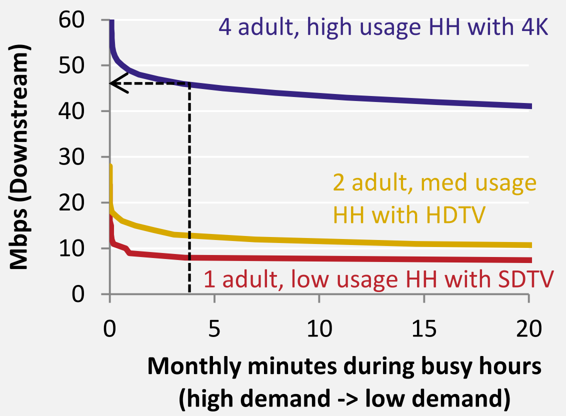

There are three important points to note regarding this approach. Firstly, while we have ranked minutes and considered the small number of busiest minutes, these minutes need not be continuous – indeed, they are likely to be widely scattered throughout the month. Further, given that much traffic (particularly web surfing) is 'bursty', required bandwidth can vary by the second.

Secondly, some high bandwidths (at the left of the graph) have very short expected durations. These expected durations are monthly averages, but in practice the relevant combination of apps may not occur at all for several months, and then (potentially) overlap for several minutes in one particular month.

Thirdly, the degree of 'magnification' in Figure 23 is high – we are considering the 20 busiest minutes out of the 7,200 minutes per month in the busy hours.

Usage profiles across households (downstream)

Having generated usage profiles as above for each of the 192 household types, we are in a position to aggregate these to understand the overall spread of demand.

Figure 24: Downstream Bandwidth demand distribution over time (4 mins excluded basis)

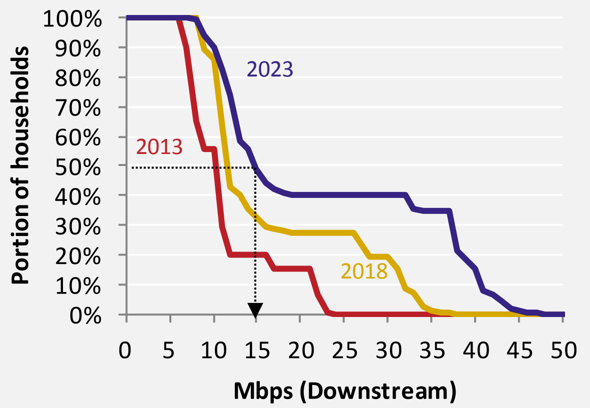

Figure 24 shows the distribution of demand over time. The curves show the percentage of households that have at least a given level of demand (on a '4 minutes excluded' basis). The arrow indicates that in 2023, 50% of households have a demand of 15 Mbps or more. (Put another way, the median demand is 15 Mbps). In that year 5% have a demand of 43 Mbps or more, and the top 1% have a '4 minutes excluded' demand of 45 Mbps.

As would be expected, these are higher levels of demand than in earlier years – the median demand in 2013, 2018 and 2023 is 10, 12 and 15 Mbps respectively (Figure 25). The demand for the top 1% of households grows from 23 to 35 to 45 Mbps from 2013 to 2023. Growth is faster for such households, which are downloading ever larger software files. By contrast, video is a more important driver of peak bandwidth for median households, and this benefits from improved compression, suppressing bandwidth requirements.

Figure 25: Downstream Bandwidth demand over time, percentile basis (4 mins excluded basis)

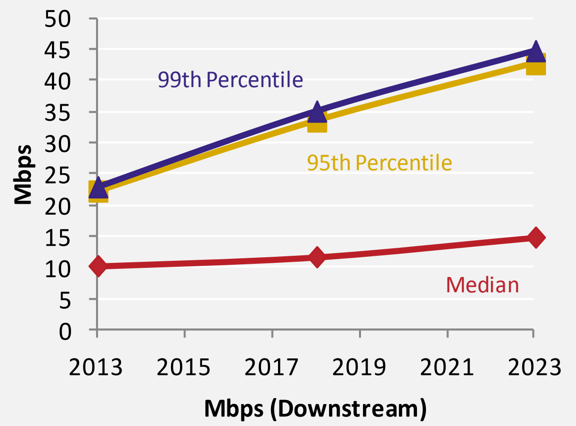

We note that an increase in median demand of 5 Mbps (from 10 to 15 Mbps) over the next ten years compares to the growth *to* 10 Mbps over the previous twenty years.[170]

A 2023 median demand of 15 Mbps may seem low, but needs to be seen in the context of the continuing benefits of video compression, and the fact that 58% of households only contain one or two people. Consider two people both surfing, both watching their own HD TV stream while each having a video call. Even this rather aggressive (and rare) use case only requires just over 14 Mbps in 2023.[171]

[170] If we date the start of the consumer internet to the 1993 launch of the Mosaic browser. At this date, websites were designed to be deliverable by 56 kbps modem (or less)

[171] 2 x HD at 2 Mbps, 2 x surfing at 5 Mbps and 2x video call at 200 kbps

Figure 26 offers another representation of the model results, as a surface plot. On the left hand axis, households are ranked from those with the highest demand (at the back left) to those with the lowest. The bandwidth demand of each household is shown on the right hand axis, from the busiest minute (at zero, on the back right wall) to the 100th busiest (at the front), out of the total 7,200 minutes in the busy hours over the month.

Figure 26: 2023 Downstream demand by minute and household

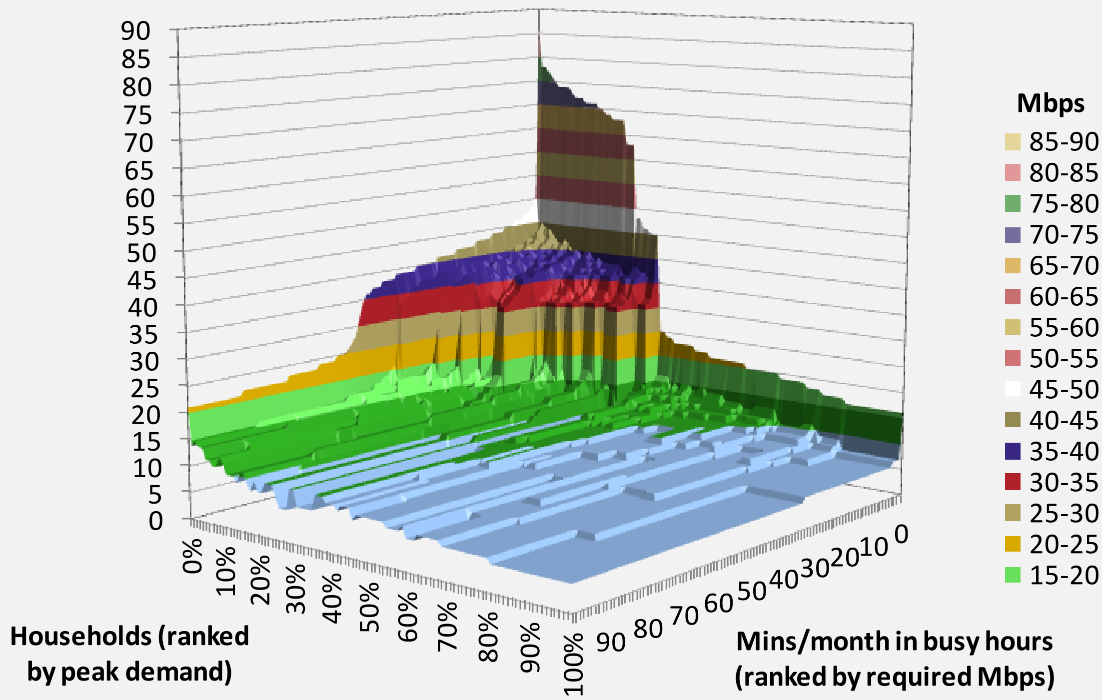

For all households there is the possibility of high demand (the 'cliff face' at the back right). However, the expected duration of that demand is short, resulting in the sharp drop off.

In the back corner, the most demanding households trigger some demand above 50 Mbps. However, such demand is neither widespread, nor sustained.

Usage profiles across households (upstream)

In addition to forecasting demand for downstream bandwidth, the model also forecasts upstream demand. As might be expected, the key applications are rather different. For example, internet TV creates much downstream demand, but drives little upstream traffic. Uploading video (to YouTube, for example) is the reverse.

Figure 27: Upstream Bandwidth demand distribution over time (4 mins excluded basis)

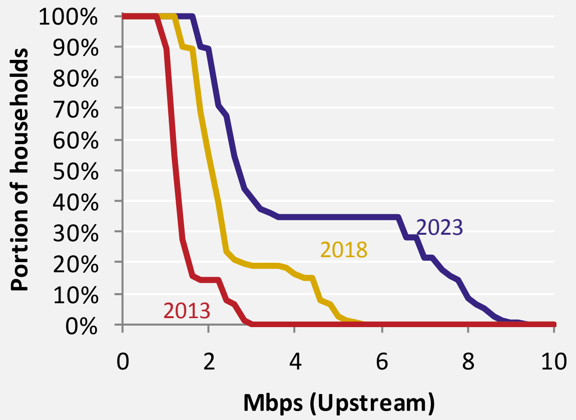

Median upstream demand grows from 1.2 to 2.6 Mbps by 2023. However, as adoption, per-user volumes and user expectations of content upload all rise, there is an increasing portion of households that require significantly more bandwidth than this.

Demand drivers and externalities

In the context of a cost-benefit analysis assessing options for market intervention, it is useful to look at the demand drivers through the prism of externalities.

Generally, where the benefits of a particular product accrue entirely to the consumer, the market operates efficiently – if suppliers can offer that product at less than the consumer is willing to pay, the market will provide the product and conversely, if the cost to provide is greater than the value to the consumer, the market will not provide, and in each case this is the 'efficient' outcome in the sense of maximizing benefit for society as a whole.

However, many products have externalities – that is, benefits that do not accrue to the purchaser. In such cases, there can be 'market failures' – the market may produce less than the optimal level of the product, because suppliers can only capture the value to the purchaser (via the sales price), and not those benefits to third parties. Governments sometimes see such externalities as a reason to intervene in markets. For instance, the benefits of education accrue not simply to the pupil and her family, but society as a whole, creating externalities that justify state funded education.

Against this background, we have sought to understand the extent to which bandwidth requirements are driven by applications with positive externalities as opposed to applications with primarily private benefits (or even negative externalities[172]).

This is not an exact science, particularly given the broad categories of application we have used for the model. For instance, accessing vital medical information or accessing a celebrity gossip site would both come under 'web surfing'. YouTube includes university lectures as well as *Gangnam Style*.

However, there are certain applications for which the benefits are likely to be predominately private (that is, are unlikely to have material positive externalities). Such applications include console game downloads; 4K TV and Bittorrent. In addition, high speed downloads of content and software, while desirable, may not bring externalities. For instance, does the difference between a ten minute download of a movie versus a twenty minute download bring benefits to anyone beyond the household in question?

By 'zeroing out' the low externality applications in the model and moderately extending download times, we can understand the extent to which the requirement for higher speeds is driven by private benefits. This in turn supports judgments as to whether society as a whole should share in the costs of providing these higher speeds, or whether the particular households desiring them should bear the full cost.

We first consider console games downloads. As we have seen, software downloads (games and operating systems) are a major driver of bandwidth requirements, and removing console game downloads has a dramatic impact (Figure 28). The 95th percentile demand drops from 43 to 23 Mbps.

Figure 28: Impact of no console games and slower downloads, 2023

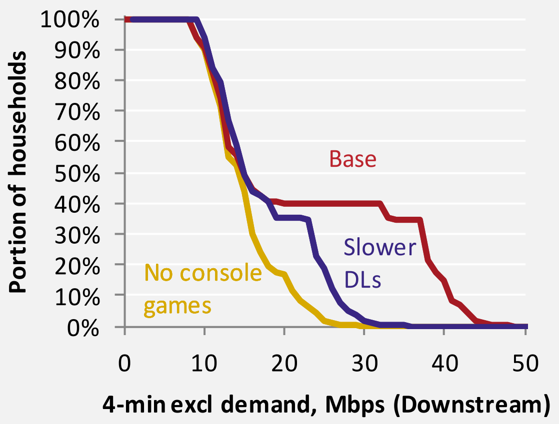

Peak bandwidth requirements are also materially affected by slowing downloads by a factor of two – for instance, extending software downloads from 10 minutes to 20. This slowing reduces 95th percentile demand from 43 to 27 Mbps.

Turning to 4K, removing this has relatively marginal impact, reducing 95th percentile demand from 43 to 42 Mbps.[173] 4K is

[172] Piracy might be an example – while the illegal downloader gets private benefit from the show in question, by taking the free download rather than paying for a legitimate copy, he deprives studios of income which might otherwise have been invested in better content for all consumers

anyway only in a certain percentage of households, and these are not necessarily the heaviest users. For instance, a 7 person household without 4K is likely to require more bandwidth than a 1 person 4K household. In addition, by 2023 the difference in bandwidth between 4K and HD is relatively moderate at 4 Mbps.

Bittorrent too has limited impact. The bandwidth used is low, particularly in busy hours, and thus the impact of excluding Bittorrent on 95th percentile demand is to reduce it from 43 to 41 Mbps.

Finally, we consider the combined impact of excluding all of these low externality applications (Figure 29). The 95th percentile demand drops to 17 Mbps (and the median demand to 11 Mbps).

Figure 29: Impact of no 'low externality' apps, 2023

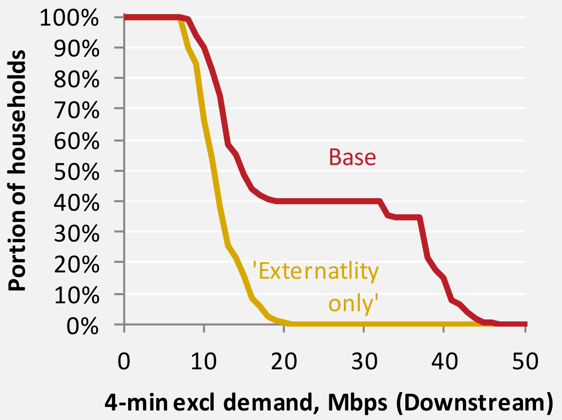

This suggests that some of the key drivers of bandwidth requirements are in fact applications with primarily private benefits.

However, in practice it may be challenging to partition usage in this way. For instance, if we wish to enable a parent to undertake CO_2-reducing home working, we may have to allow for the possibility that a teenager in the next room is streaming *The Desolation of Smaug* in 4K. Put another way, if a household is bandwidth constrained, there is no guarantee that the dropped application is the one with only private benefits.

Sensitivities to model results

As with any model, this one is dependent on a large number of assumptions. Particularly for a model looking out ten years, as this does, there is considerable uncertainty attached to these assumptions. Reasonable people may legitimately disagree with some or all of our inputs.

Therefore we have undertaken a sensitivity analysis to illuminate the impact of varying certain key assumptions. Some of these represent different possible real-life outcomes (for instance, the rate at which video compression improves), others represent

[173] Note that even if the 'before' 95th percentile household was using 4K TV, and removing 4K substantially reduced that household's demand, this need not have a significant impact on 95th percentile demand. If the previous 94th percentile household did not use 4K TV, then its demand is unchanged, and it might 'step in' to become the setter 95th percentile demand

different conceptual approaches (for instance, excluding more or fewer minutes when assessing bandwidth demand).

Demand criteria

Taking first the issue of the number of minutes to exclude, Figure 30 shows the impact of tightening or loosening the '4 minutes excluded' threshold. If it is tightened to 1 minute, then the required median bandwidth rises from 15 to 17 Mbps, and the 95th percentile demand from 44 to 45 Mbps. Conversely, if the criterion is loosened to 30 minutes (equivalent to one degraded minute per household per day), the median required bandwidth drops to 11 Mbps and the 95th percentile demand to 32 Mbps.

Figure 30: 2023 Downstream Bandwidth demand distribution (various mins excluded basis)

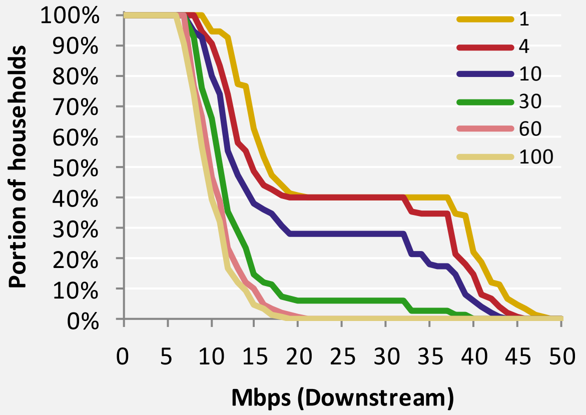

Usage concentration

Our base 'High/Medium/Low' assumption is that 20% of individuals have twice the average usage, 40% have average usage and 40% have half the average. To explore the impact of usage being more skewed than this, we have considered a scenario where 10% of the users have triple the average usage (and the number with average usage expands to 50%). Figure 31 shows the results. At the top end, 95th percentile demand increases from 42.7 to 43.2 Mbps.

Figure 31: 2023 Downstream Bandwidth demand distribution (more concentrated usage)

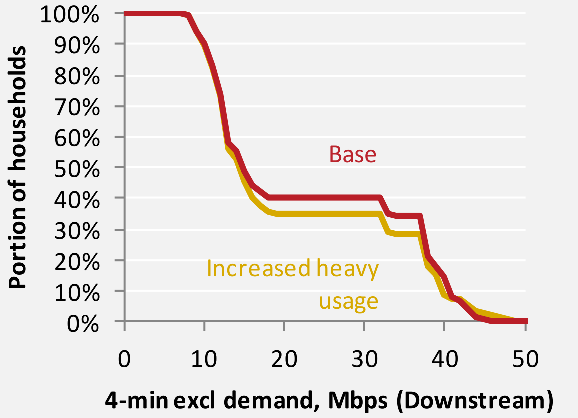

This change is relatively moderate since increased usage does not translate very directly into increased bandwidth requirements. For example, if a given household's 4 minute bandwidth requirement already included an Internet TV stream, the fact that there is greater consumption of Internet TV at other times makes no difference to the 4-minute-excluded demand.

Future scenarios (downstream)

We now turn to the impact of varying key assumptions about the future, first looking at downstream demand. Figure 32 shows the impact of varying assumptions related to downloads and surfing. As would be expected, if user expectations for download or page load times are more ambitious than that shown in the base case, then required bandwidth rises.

Figure 32: 2023 Downstream download and surfing scenarios (4 mins excluded basis)

Assumption	Base	Sensitivity
Video content DL (as portion of real time)	25%	10%
Web page load time	3 secs	2 secs

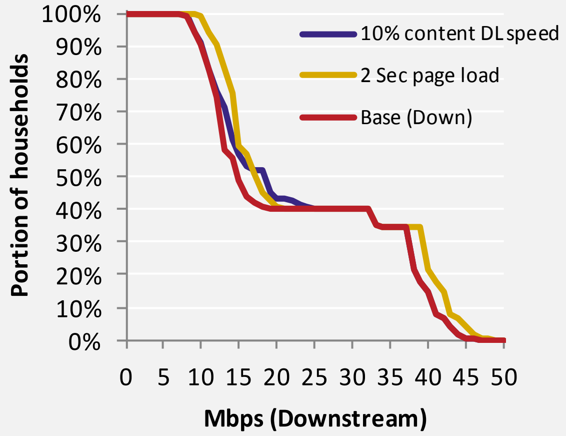

Since web usage is widespread and extensive, more rapid page loads increase the bandwidth requirements for almost all households, by approximately 2 Mbps. More rapid video content download speed increases the median speed requirement from 15 to 18 Mbps. However it has minimal impact on 95th percentile demand. This is because such demand is driven significantly by console game downloads. Because both console and content downloads are relatively infrequent, the probability of overlap of the two is very low, and hence the impact of changes to content download speed on this high end demand is insignificant.

Figure 33 looks specifically at expectations of software download times (primarily console games). As we have noted, as with other download time assumptions, our base case assumption of 10 minutes is arbitrary, and reducing this expectation increases bandwidth requirements. For instance, if it is reduced to 6 minutes, then 95th percentile demand rises to 64 Mbps (from 43 Mbps). At 4 minutes it rises to 90 Mbps, and tightening this further would quickly take demand in these most demanding households above 100 Mbps.

Figure 33: 2023 Downstream software DL expectation scenarios (4 mins excluded basis)

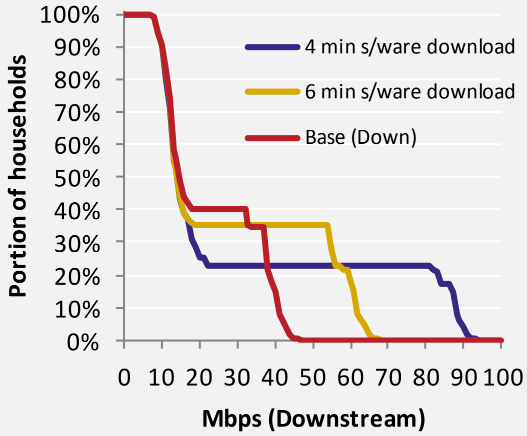

Note however that the downloading of large software files is a comparatively rare event. Consequently, for some households tightening the download time expectation means that downloading is reduced to less than 4 minutes per month, and so is excluded from the demand picture shown here. (Hence the reduction in demand for households in the 23-35% bracket as expectations tighten from 6 to 4 minutes).

Figure 34 considers scenarios related to usage overlap and internet TV. By reducing the number of busy hours per day from 4 to 2, while holding the amount of usage in the busy hours constant (at 50%) we can see the potential for increased concentration to drive greater bandwidth demand. In fact, the impact is relatively moderate - 95th percentile demand increases from 43 to 46 Mbps. This is because even at four hours, there is already a large amount of overlapping usage.

Figure 34: 2023 Downstream Internet TV and busy hour scenarios (4 mins excluded basis)

Assumption	Base	Sensitivity
Busy hours (containing 50% of traffic)	4	2
Initial 4K bandwidth (Mbps)	16	30
Annual improvement in video compression	9%	6%

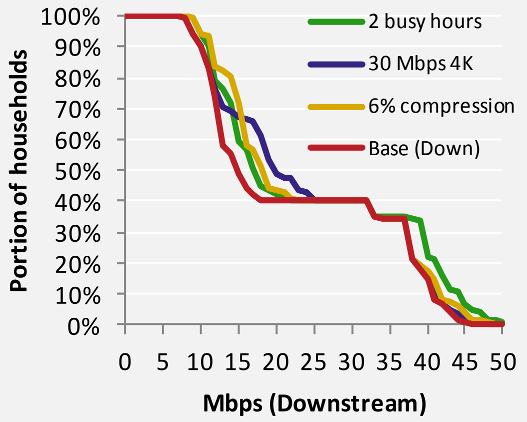

An increase in starting (2013) bandwidth for 4K TV from 16 to 30 Mbps again has only moderate impact – in part because compression narrows the gap between these two starting assumptions by 2023. The impact is less significant in downloading households, since in such homes it is more likely to be software downloads which sets peak demand rather than Internet TV. However, in the median household demand increases from 15 to 20 Mbps.

A change to slower improvements in bandwidth compression affects all households. Video compression drives required bandwidth for Internet TV, for content downloads, for video calls and various other applications and so has broad impact. A 6% annual improvement (rather than 9%) increases 95th percentile demand to 45 Mbps.

Above we have considered the impact of changing various assumptions individually. We have also considered the impact of changing them together, in an 'aggressive' scenario that combines a range of such assumptions (Figure 35). Note that this scenario is not an ‘upper bound’ of potential demand – it is possible to imagine users being even more impatient for software downloads, for example. (Below we explore a more aggressive scenario).

Figure 35: 2023 Aggressive scenario (4 mins excluded basis)

Assumption	Base	Aggressive
Mobile OS DL time	15 mins	10 mins
Video content DL (as portion of real time)	25%	10%
Web page load time	3 secs	2 secs
Software DL time	10 mins	5 mins
4K TV penetration	43%	60%
Busy hours	4	2
Video compr’n CAGR	9%	6%
Initial 4K bandwidth	16 Mbps	30 Mbps

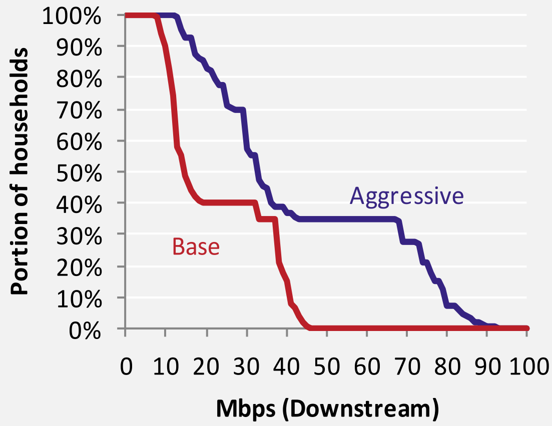

In this scenario median ‘4 minutes excluded’ demand increases from 15 to 33 Mbps, and 95th percentile demand from 43 to 84 Mbps.

Another way to consider sensitivities in the model is ‘what do you have to believe to see major demand above 100 Mbps’. To explore this we have looked at a ‘Very aggressive’ scenario. Note that we do not regard this scenario as very plausible, since it involves very aggressive assumptions on many dimensions – it is offered more as a thought experiment.

Figure 36: 2023 Very aggressive scenario (4 mins excluded basis)

Assumption	Base	Very Aggressive
As for aggressive scenario, plus:		
Internet TV / total TV	25%	50%
Web page load time	3 secs	1 sec
Software DL time	10 mins	2.5 mins
Software DL/year	1.5	3
YouTube minutes	1000	1800
Web usage growth	3%	5%

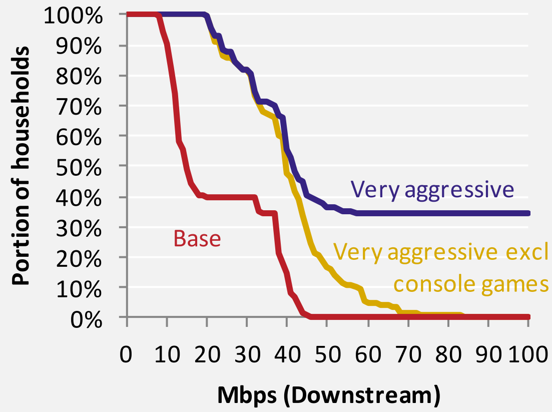

This very aggressive case involves 5.4 hours per person per day of conscious internet usage in the home, which is clearly very substantial (and more for heavier users).[174] However, even so only 35% of households have bandwidth requirements exceeding 100 Mbps. Moreover, this level of need is driven entirely by games downloads – if these are excluded, then demand drops dramatically, with 95th percentile demand at 62 Mbps.

As we discussed above,[175] 'faster than real time' applications such as console game downloads have a distinct role in setting bandwidth requirements. In households limited to real time applications, bandwidth requirements are likely to be more limited, almost regardless of usage levels.

As a consequence, we have not been able to construct a credible scenario in the model which sees demand above 100 Mbps other than driven by console game downloads.

Downsides to the base case

In the discussion of sensitivities and scenarios above, we have focused on sensitivities that would drive increased demand. However, we have done so because we are aware that our base case forecast is lower than some more casual forecasts, *not* because we believe these upside sensitivities are more likely than downside sensitivities. Indeed, we believe that our model is, in a number of regards, highly conservative. Outcomes with lower demand than we forecast are entirely plausible. For instance:

- In considering web usage, we have not taken account of the substantial shift to browsing on a mobile device, which will appreciably reduce effective page weights
- We have assumed that 4K TV is a success, with widespread adoption of sets and availability of on-demand 4K content. However several commentators are sceptical of the benefits and prospects of 4K TV[176]
- We have factored in both ongoing substantial games downloads, and increasing cloud-based streamed gaming. The latter could substitute for the former, greatly reducing our forecast bandwidth requirements
- We assume substantial growth in both legal and P2P content downloads, as well as dramatic growth for streamed internet TV. There is evidence from other markets

---

[174] See page 64 for a discussion of time available for internet use
[175] Page 49
[176] See for example, CNET, *Why Ultra HD 4K TVs are still stupid*, 29 January 2013; Roy Furchgott, "Why You Don't Need a 4K TV", *New York Times*, 17 June 2013; Bob Zitter [HBO CTO] quoted in EOSHD, *3D a proven failure, 4K unlikely to succeed – HBO*, 26 March 2013; The Economist, *After 3D, here comes 4K*, 9 March 2013; Fox News, *4K TV: Fad or Fantastic*, 30 April 2014

that streamed TV in fact acts as a substitute for content downloads

- For downloads we have assumed our calculated average speed requirement (for instance, to receive a movie in twice real time) is made available throughout the download. In practice, speeds could be lower at certain times, for instance to accommodate some other spike of demand, as long as there was spare capacity later for the download to 'catch up'
- We have not factored in the human limits to multitasking – for instance, we have not assumed that having two tasks in process makes an additional simultaneous task any less likely

Given the above, we are comfortable that our base case strikes a reasonable balance between upside and downside risk.

Future scenarios (upstream)

We have also considered sensitivities for upstream demand (Figure 37). Increased usage concentration (2 busy hours) makes relatively little difference. This is because upload requirements are not particularly driven by overlapping usage – rather it is the spikes in demand due to one particular application, video and content uploads, which are key.

Figure 37: 2023 Upstream scenarios (4 mins excluded basis)

Assumption	Base	Sensitivity
Busy hours (containing 50% of traffic)	4	2
Video compression CAGR	9%	6%
Video upload (% of real time)	50%	25%
Aggressive		[All of the above]

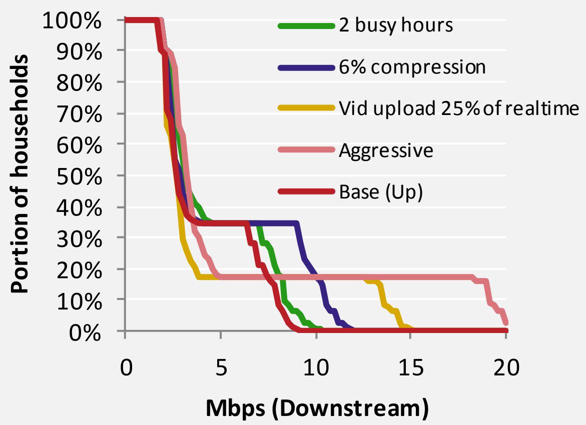

If user expectations of such uploads are that they will take place in 25% of realtime (rather than 50%), this has a dramatic impact. The 95th percentile households will require 14.8 Mbps or more, albeit briefly. (For households in the 16-36% tier, these faster uploads drop total uploading time below the 4 minutes per month threshold, and so this demand is excluded in the chart above).

Once again, compression is an important assumption, since it drives the size of video files being uploaded, and by extension the required bandwidth.

We have combined these sensitivities to create an aggressive upstream scenario. This sees 95th percentile demand rise to 19.9 Mbps.

Comparison to other forecasts

The base case of the model suggests that for all but 1% of households, 45 Mbps downstream is likely to be sufficient even in 2023 (on a 4 minutes excluded basis). Even under the combined impact of a range of more aggressive assumptions, this figure rises only to 90 Mbps.

These figures are lower than those sometimes informally offered for likely future bandwidth demand (though they are consistent with empirical evidence suggesting that even at peak demand, few users currently saturate their broadband connections).[177] Our results are a mathematical consequence of our input assumptions (which, as noted, are open to debate), but at a high level we believe the reasons that we differ from some other forecasts are as follows:

We have not presumed capacity and demand will move together

In the past, internet access has been capacity constrained, with clear anticipated applications that were not yet possible.[178] As a consequence, increases in capacity led quickly to increases in demand. To take a simple example, Internet TV has been being expected since the days of 28.8 kbps modems.[179] As capacity grew, there were immediately available types of content to make use of it (and often pools of pre-existing content). Basic broadband enabled the web with pictures and audio, faster broadband enabled YouTube and today's broadband has enabled iPlayer and Netflix.

However this past strong relationship between capacity and demand may now be breaking down, not least because (as the FTTH Council says) there is "no really compelling application yet" for faster speeds, a stark contrast to earlier steps in the development of broadband.[180]

In allowing for the possibility of a divergence of capacity and demand, we are implicitly rejecting the view that 'build it and they

[177] Sarthak Grover et al, *Peeking Behind the NAT: An Empirical Study of Home Networks*, October 2013
[178] Certainly for some households this remains so today, as the model suggests
[179] Edmund DeJesus, "How the Internet will Replace Broadcasting", *Byte*, February 1996
[180] FTTH Council Europe, *Creating a brighter future*, 20 February 2013

will come'. As we have noted, the model does not (and could not) treat 'unknown unknowns', bandwidth uses that we have not yet imagined but which might be enabled by the availability of higher speeds. These thus represent upsides to our forecast.

Growth is likely to slow as we approach 'human limits'

Another reason why past growth is not necessarily a good guide to future growth is that on some dimensions we are approaching 'human limits'. People only have so many hours in a day, their ability to multitask is finite and their visual acuity is limited. In the past, substantial growth in demand operated well within these limits, but they are likely to begin to be constraints. For instance, it is hard to see what visual benefit a TV generation beyond 4K might bring.[181] (By contrast to the clear benefits of moving from the rudimentary video of a decade ago to today's Internet TV). In this context, it is notable that the bandwidth of the human eye is estimated at 10 Mbps.[182]

Figure 38: Time available to be online at home [183]

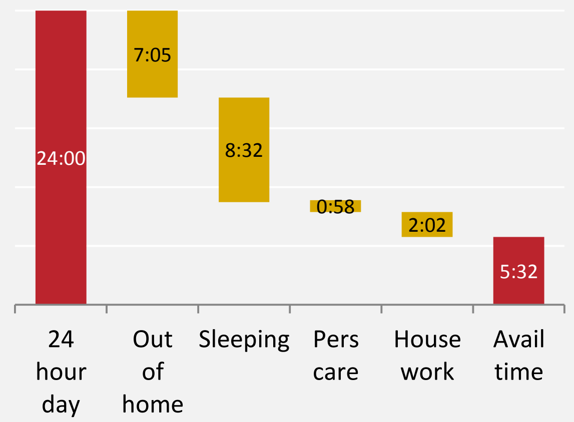

Equally, while the model anticipates 'conscious usage' of the internet[184] in the home rising to over 3.8 hours per day per individual in 2023 from 1.8 today, this is a slower rate of growth than in the past.[185] This slower rate of growth seems inevitable, in that with 3.8 hours of domestic internet usage, Australians will be spending a significant majority of available hours online.

The average person spends just over 7 hours out of the home per day, and a further 8½ hours asleep. Of their conscious time in the home, almost an hour is spent on personal hygiene and health care, and two hours on house work.[186] While internet use is not impossible simultaneous with such activities – internet radio is compatible for example – it seems likely than the great majority of internet usage will take place in the remaining 5½ hours (time

[181] Some are doubtful of the benefits even of 4K – Forbes has called it "not only an extravagance, but an invisible extravagance to anyone with regular visual acuity". Sharif Sakr, "How Long Before A 4K TV Becomes A Realistic Purchase? Give It Two Weeks", *Forbes*, 9 April 2013

[182] University of Pennsylvania, "Penn Researchers Calculate How Much the Eye Tells the Brain" [Press release], 26 July 2006

[183] Communications Chambers analysis of ABS, *How Australians Use Their Time*, 2006

[184] Excluding background activities such as file downloads. Note that the 4.5 hours is the sum of individual uses. Since these may be multitasked (for example, surfing and internet TV), the elapsed time will be less

[185] Note that the 3.8 hours figure understates total conscious usage anticipated by the model, since it excludes applications such as streaming radio, online gaming and home working via VPN, all of which are grouped together (with other applications) in our 'low bandwidth' category

[186] Food preparation & clean-up and grounds and animal care are the majority of this

which also has to accommodate a wide range of other activities, such as eating, childcare and so on). Since we are forecasting 3.8 hours of conscious internet usage, we are in effect assuming that a significant majority of time available in the home for internet use *will* in fact be spent online.

Moreover, our 3.8 hours is an average figure, but the model works on the basis that 20% of individuals are heavy users who have double this level of usage. (These are the users who, generally, drive 95th percentile demand). These users have 7.6 hours of conscious usage in the home per day, which suggests substantial multitasking.

However, there are certain dimensions where we are some way from human limits, in particular impatience. As we have seen, user expectations of speed of content and software downloads is an important sensitivity.[187] Those unwilling to wait more than seconds for very large files to download may demand very high bandwidths.

We are focused on bandwidth not traffic

Perhaps the most familiar internet forecasts are the Cisco VNI figures. These however are forecasts of total traffic, not bandwidth per home. As a result they include growth due to increased internet adoption, both fixed and mobile. Particularly at an international level (which are the figures most often quoted), these are very important factors. Even allowing for the impact of adoption, traffic growth is very different from bandwidth growth, as we have noted.

That said, the details of the Cisco Australian traffic figures are instructive. Between 2012 and 2017, they forecast that consumer fixed internet traffic will grow by 19% annually. (Our model predicts a 26% CAGR, 2013-17). However, Cisco's growth is almost entirely due to an increase in video, which they see growing at 24% annually from an already high base. They forecast a growth rate of 7.4% for their 'web and other' category, and see filesharing traffic *falling* by 0.6% annually. (These are their three key categories of traffic).

Figure 39: Cisco VNI forecast of Australian consumer fixed internet traffic[188]

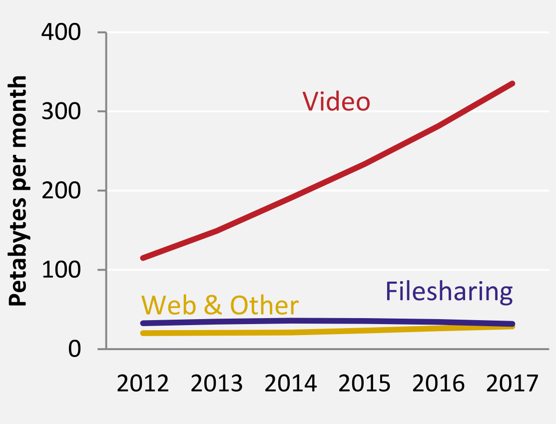

This variation in growth across these categories is particularly relevant in the context of understanding bandwidth. Video is a very

[187] See page 53
[188] Derived from *Cisco VNI Forecast widget*, set for Australia consumer fixed internet traffic [accessed 14 April 2014].

significant driver of traffic, but (in common with all streaming applications) is relatively less important in bandwidth terms. This is because its traffic (and traffic growth) derives substantially from the long and growing periods for which it is used. However, this needn't drive bandwidth growth. If a household's peak bandwidth need is set by 4K Internet TV usage, for example, its need will be the same whether there is one hour of usage per month or ten.

Many bandwidth forecasts rely on 'informal' views of application bandwidths and their combination

Much of the discussion of forecast bandwidth needs has been based on what might be called 'informal' views of the bandwidth requirements of particular applications. For instance, we have noted that WIK's recent work on high speed broadband in Germany takes as an input assumption that HD videocommunication requires 25 Mbps.[189] This unsourced view compares to Skype's requirement of 1.5 Mbps for HD video calls.[190] (Even group video with 7 parties – presumably a *very* rare event at home – only requires 8 Mbps according to Skype). Clearly such radical differences in assumptions will lead to very different bandwidth forecasts. We note also that bandwidth projections generally take no account of improvements in video compression.

Many bandwidth forecasts take no account of the duration or probability of the highest demand scenarios

Not only are the bandwidths of the individual applications in many usage projections open to question, so too are the way in which these are combined. Scenarios are often painted that in reality would require seven or eight people in a household to be online simultaneously.[191] Such scenarios are not impossible, but our modelling suggests they are extremely rare – it is for this reason that we think that it is vital to consider the expected duration of any particular demand scenario. More generally, we find discussion of bandwidth needs rarely takes into account that a significant majority of households contain only one or two persons.

Further, unlike some other forecasts, we have not worked on the basis that all demand (no matter how brief) must be met – instead we have taken our 'x minutes excluded' basis, albeit with a tight threshold of just one minute per week. This inevitably leads us to lower headline numbers for future demand, although (for the reasons set out above) we believe this is appropriate.

---

[189] WIK, Market potential for high-speed broadband connections in Germany in the year 2025, 15 January 2013

[190] Skype, How much bandwidth does Skype need?, accessed September 2013

[191] See for instance Trent Williams, NBN Co Community Engagement : Working with Communities and Councils, 2011 and Ioannis Tomkos, Techno-economic Evaluation of Next Generation FTTx Access Network deployments, 11 February 2011

7. Conclusions

As with any model, this one has limits. Precision must be traded off against complexity, data is often best-available rather than perfect, and so on. Also, as with any forecast, it includes a range of assumptions that are the best guess of the authors as of today, but which may prove inaccurate.

Moreover, it is a forecast of a highly dynamic system (the internet) which will continue to be in flux. While the model anticipates many changes in usage (and uses broad categories of applications to avoid the need to precisely predict specific applications), by definition it cannot predict 'unknown unknowns'. Highly demanding applications may appear which are outside any of the categories in the model. These represent upsides to our forecasts, and so it may (subject to cost) be sensible to 'overbuild' relative to these forecasts to allow headroom for such potential applications.

As we have noted, the model is also dependent on some highly arbitrary assumptions about consumer expectations of download speed for software and content. Others may (legitimately) take a different view on these inputs.

For these reasons, the precise outputs from the model are open to debate and we see our modelling exercise only as a contribution to a meaningful discussion about bandwidth needs, rather than its conclusion. However, we feel strongly that this discussion will be better founded if it is based on:

- Rigorous analysis of the probability of app stacks
- Reference to the actual bandwidth requirements of individual apps (and their increase or decrease over time) rather than loose estimates
- An understanding of actual household demographics rather than on a notional 'typical' household (which is frequently anything but)
- An understanding of the duration of peak demands

Finally, we note the limits of the model's scope – it is a forecast of bandwidth requirements. However, required bandwidth is only one input to a consumer purchase decision, and this study does not look at other reasons as to why consumers may wish to take up a higher (or lower) bandwidth service or what products ISPs will seek to offer consumers over time.

8. Assumption summary

				2013	2023
Busy hours traffic concentration					
Busy hours per day			4.00		
Portion of traffic in busy hours					
General			50%		
BitTorrent			30%		
Child:adult usage in busy hours			50%		
Household usage variation					
	Household usage type				
Share of	*Low*	*Med*	*High*		
Households	40%	40%	20%		
Usage	20%	40%	40%		
Video bit rates (Mbps)					
SD				2	1
HD				5	2
4K				16	6
4K TV penetration				0%	43%
Video compression improvement (per year)			9%		
Video usage (Mins/adult/day)					
Internet TV				2.7	46.5
YouTube etc.				17.7	33.3
Content downloads					
Hours of video downloaded per month				0.43	1.12
DL time as fraction of real time				50%	25%
Portion of users downloading content				10%	40%
Bittorrent					
Hours of video downloaded per month				16.7	52.8
Users				30%	40%
Mobile OS downloads					
Download time expectation (mins)				30	15
OS Size (MB)				700	1,140
Software downloads					
Download time expectation (mins)				60	10
File Size (GB)				7.0	23.7
Pre- & post-loading benefit				0%	90%
Using households				20%	40%
Downloads per month (high usage HH)				0.4	0.4
Video calling (SD)					
Usage (mins/adult/month)				40	400
Bandwidth (Mbps)				0.50	0.19
Compression improvement			9%		
Video uploads					
Using households				20%	40%
Video mins/indiv/month uploaded				5.00	20.64
Uploaded bitrate				2.00	3.15
Upload time as fraction of real time				100%	50%
Web usage					
Usage (hours/adult/day)				1.45	1.95
Average page weight (MB)				1.70	5.87
Max 'above the fold' page weight (MB)				2.00	2.00
Time on page (seconds)				75	91

9. Glossary

µTP	A protocol used by the BitTorrent file sharing platform
4G / LTE	A standard for wireless communications of high-speed data for mobile devices
4K TV	Standard of video with a much higher resolution than high definition TV today, typically 4,096 horizontal pixels (hence 4K). 4K TV is often used interchangeable with Ultra High-Definition (UHD) TV, but is actually a slightly higher resolution standard (UHD TVs typically have resolution of 3,840 x 2,160)
'Above the fold'	The portion of a web page that is visible in a browser when the page first loads (contrasted with 'Below the fold')
Bandwidth	A measure of the actual or potential rate of transfer of internet data, expressed as a quantum of data per second
CAGR	The compound annual growth rate, representing year-on-year growth over a number of years
Cloud computing	Applications and services offered over the internet. Cloud applications are typically opened through a web browser rather than running installed on a computer.
DVR	A digital video recorder, a video recording device which records onto a hard disk drive (rather than a disc or tape)
'Fine nines'	A telecoms term indicating 99.999% uptime, referring to the high availability of a service (when downtime is less than 5.26 minutes per year)
FTTB	Fibre to the building, a type of network architecture which uses optical fibre cable goes to a point on a shared property
GB	Gigabyte, a unit of data which equals 8 Gigabits (Gb) and 1024 Megabytes (MB)
Gb	Gigabit, a unit of data which $1/8^{th}$ of a Gigabyte (GB) and equals 1,024 Megabits (Mb)
HD	High Definition, a higher resolution of video than Standard Definition. HD TV can be transmitted in various formats, the most popular including 720p, 1080p and 1080i

Internet TV	Professional television / video content distributed over the internet (rather than broadcast networks). We use the term to include both consumption via TV sets and via other devices
ISP	Internet service provider, a company which supplies internet connectivity to a home or business
KB	Kilobyte, a unit of data which equals 8 Kilobits (Kb) and 1,024 Kilobytes (KB)
Kb	Kilobit, a unit of data which 1/8th of a Kilobyte (KB) or 1,024 bits (b)
Kbps	A measure of bandwidth in kilobits (Kb) per second
MB	Megabyte, a unit of data which equals 8 Megabits (Mb), 1,024 Kilobytes (KB) and 8,192 Kilobits (Kb)
Mb	Megabit, a unit of data which is 1/8th of a Megabyte (MB) or 1,024 Kilobits (KB)
Mbps	A measure of bandwidth in megabits (Mb) per second
M2M	Machine to machine, technology which allow similar devices to exchange information and perform actions without the manual assistance of humans
OS	Operating system, the software that manages hardware and other software and application on a device or computer. Examples include Apple iOS, Google Android, Microsoft Windows and Linux
P2P	Peer to peer, used to describe applications in which users can use the Internet to exchange files with each other directly or through a mediating server (e.g. BitTorrent)
PB	Petabyte, a unit of data which is equal to 1,000 terabytes (TB) and 1m gigabytes (GB)
Primary applications	By our definition, applications that are primarily used 'one at a time' by a given individual (though they may be used in parallel with non-primary apps).
SD	Standard definition, a video standard with resolution below that of High Definition (HD)
Secondary applications	By our definition, applications that are amenable to multitasking - for instance, launching a movie download and then continuing with other activities, such as web use

TB	Terabyte, a unit of data which is equal to 1,000 gigabytes (GB) and 1m megabytes (MB)
Torrent / BitTorrent	A protocol supporting peer-to-peer file sharing that is used to distribution data over the internet
Traffic	The amount of data sent and received over the internet
'x minutes excluded'	Our metric which measures the bandwidth that would be necessary to serve all but the x busiest minutes in the month for a given household