Robert Kenny & Tom Broughton

Domestic demand for bandwidth

An approach to forecasting requirements for the period 2013-2023

5 November 2013

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.

BSG foreword

At the start of 2013, the Broadband Stakeholder Group (BSG) announced its intention to focus upon and inform the demand side element of the overall broadband policy debate.

Following a lengthy and proper focus on the costs and capabilities of different technologies to support the delivery of government broadband speed and coverage targets, we believe the overall policy debate needs to pay as much attention to the demand side. Or in other words, to consider uptake and usage rates for broadband, what people use (and don't use) their broadband for, what they want out of their broadband connection and what the overall socio-economic benefits are of that usage.

We believe a better understanding of these issues is critical to determining whether UK broadband is on the right trajectory or not, including the relative emphasis given to factors such as coverage, take-up and speed. As the government commences its development of a new Digital Communications Infrastructure Strategy for 2015-2025, this demand question should be a central foundation for that analysis.

Accordingly, the BSG commissioned this report from Communications Chambers as an input to the demand-side debate. The scope of this report does not cover all demand-side issues but explains the methodology and results of a model Communications Chambers has developed that seeks to forecast UK domestic demand for broadband capacity for the period 2013-2023.

We were motivated to commission such a study for a number of reasons. Firstly, despite the global interest in whether broadband infrastructure is meeting demand and if it will stay ahead of that demand curve going forward there is arguably a lack of material on how one might model and measure that demand. Furthermore, we were interested in developing a technology-neutral approach to forecasting demand, rooted in the behaviours of consumers and the services they want to access over broadband.

Consequently, the approach Communications Chambers has developed:

- Is fully transparent about the methodology deployed and assumptions made
- Is anchored in the speeds required by different types of applications
- Accounts for likely changes in speed requirements over time and whether they increase or decrease
- Reflects the variations in demand across different household types in the UK
- Builds a quantified view regarding the probable and likely duration of different application stacks (i.e. when will people use applications simultaneously and for how long)
- Provides a picture of duration of peak demand in a household to consider how many households will need a certain amount of bandwidth for a certain amount of time
- Focuses on consumer use in the household – whilst applications typical to basic home-working are included, more specific business use cases are not covered

Taking this approach, Communications Chambers has combined the usage profiles of various applications into the usage of profiles of individuals and from this developed these individual profiles into household profiles. In all, 156 household profiles are modelled, based on their demographics (that is, the number of adults and children present); their intensity of use; and their TV type. These 156 household profiles are then combined into a picture of national demand, demonstrating the likely difference in demand across household types.

For example, in a single person household with SDTV, the model predicts that in 2023 the broadband connection is idle for most of the time with several hours per month requiring 5 Mbps and shorter periods within the range of 8-10 Mbps. In contrast, a high use, 4 adult household with a 4K TV sees appreciable usage almost constantly during the busy hours, with approximately 90 minutes of demand of 25 Mbps or more per month.

Looking across all households, the model indicates that the median household will require bandwidth of 19 Mbps by 2023, whilst the top 1% of high usage households will have demand of 35-39 Mbps.

At first glance these numbers might seem low compared with other speeds commonly cited regarding future capacity needs. However, we would highlight a number of factors the report raises which have a bearing on this:

- It is important to remember that 64% of UK households are one or two people and therefore there is a natural limit to the online activities of these majority households. For example, even if two people are each watching their own HDTV stream, each surfing the web and each having a video call all simultaneously then the total bandwidth for this use case is 15 Mbps in 2023.
- Another important consideration is compression. The bit rate for a given quality of video has fallen steadily over time and will continue to do so due to improvements in compression techniques. As such the model assumes an annual improvement in compression of 9% for SD, HD and 4K TV.
- Furthermore, we must be careful not to confuse traffic volume forecasts (which continue to predict rapid and continued growth) with what might be expected for bit rate. Video is a considerable driver of traffic and broadband usage but is relatively less important in bit rate terms. Existing access networks could potentially absorb greater amounts of traffic without upgrade. Hence an increase in traffic does not necessarily equal a correlated increase in maximum bit rate requirement.

We should also be explicit that Communications Chambers has, in its model and report, presented a middle case falling between an evolution of today's consumer expectations of the performance of applications and a perfect world where all applications would be instantaneous. The results are also presented on a '4 minutes excluded monthly' basis. What this means is incorporating all required demands except for the 4 busiest minutes in the month.

Clearly this approach is ripe for interrogation and debate and as the report makes clear, reducing the excluded minutes pushes up the requirement. For example for a 4 adult, high usage household with a 4K TV, reducing the excluded minutes from 4 to zero would push the bit rate requirement in 2023 up from 38 Mbps to over 50 Mbps.

In that vein, the report also highlights a number of sensitivities to the model results which could change anticipated requirements. These factors include changing user expectations for factors such as download speeds and notably, reducing the time one would expect a software download, such as a console game, and upload of files to take. For example, in significantly reducing the base case assumption of 10 minutes waiting time for a console game to 2.5 minutes, then 16% of households require 83 Mbps. Reducing the waiting time further would quickly take demand over 100 Mbps for those households.

Accordingly what we are keen to emphasise in publishing the report, is that we are not presenting or endorsing an absolute "truth" on speed needs over the next decade. As with any report looking forward, one cannot predict the future with exact certainty.

Furthermore we should be clear about what is not within the scope of this study. It does not look at the ways in which one might seek to build demand for higher speed networks or services, nor assess potential innovation of applications over superfast and ultrafast networks.

However, what this report does do is set out a robust, clear and transparent explanation of its approach to forecasting demand. In line with the BSG's overall objective to put evidence-based information into the public domain to improve the quality of the public policy debate about broadband, we believe this report fills a significant evidence gap. As is often the case in developing a model to measure potential future behaviour, the value is as much in the rigorous analysis of the problem as in the output.

Ofcom's 2013 Infrastructure Report reports availability of superfast broadband services (defined as delivering a speed of at least 30 Mbps) to 73% of UK premises, 22% of fixed line internet subscribers adopting those superfast services and increasing 4G coverage in the UK. Looking at the impact of connection speed on data use, it also highlights the threshold at which data consumption plateaus at 10 Mbps, up from 8 Mbps in 2012.

In the context of these findings, the BSG believes that this report provides new insight and evidence to inform a variety of key broadband policy questions. In publishing the full underlying model alongside this report, the BSG wants to instigate an informed and evidence-led discussion about the future demand for bandwidth and what this means for UK broadband policy overall, including the issues of:

- Will anticipated demand be met by current and anticipated infrastructure provision from both the private and public sector?

- Will that infrastructure provision meet demand across all geographies, such as the final 10%, over time?
- What does this picture tell us about thresholds for basic broadband provision as the government seeks to build on its 2 Mbps commitment and support the provision of superfast services more broadly?
- What does this mean in the context of encouraging more people to take-up and use the broadband services available to them and the correct balance between this public policy objective and the drive for increased speed and capacity?
- How are consumer expectations likely to change in terms of what they want to use their broadband for, where they want to access connectivity and how will this impact on network provision?
- To what extent will network capacity respond to demand and to what extent might increased network capacity drive demand?
- What will the future picture of demand mean for policy making beyond the parameters of the Digital 2020 targets?

As the BSG works with industry and wider stakeholders on what factors the government needs to consider as it develops its longer-term digital strategy, this debate about the interrelation between network infrastructure and what people want and need to do over these networks becomes ever more important to define and understand.

This report brings significant new insights to this core policy question and the BSG looks forward to taking this debate forward.

Contents

1. Executive Summary

This paper sets out the methodology and results of a model[1] that seeks to forecast UK domestic demand for broadband capacity. As far as we are aware, it is the first such model to be put into the public domain. It has been developed by Communications Chambers, and commissioned by the Broadband Stakeholder Group (BSG) with support from BSkyB, BT, Ofcom, TalkTalk, Three and Vodafone.

Note that 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).

Objectives

Given the worldwide interest in increasing broadband speeds, it is perhaps surprising that there is so little methodical work on what consumers' requirements[2] 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 64%[3] of UK households contain only one or two people. 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 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.

[1] The model itself is available from the BSG website: http://www.broadbanduk.org/forecastingdomesticdemand
[2] We use requirements in the technical sense of bandwidth consumed, rather than speed that a customer might be willing to pay for (which could be more or less)
[3] ONS, *Families and Households 2012*, 1 November 2012

- 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 14 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

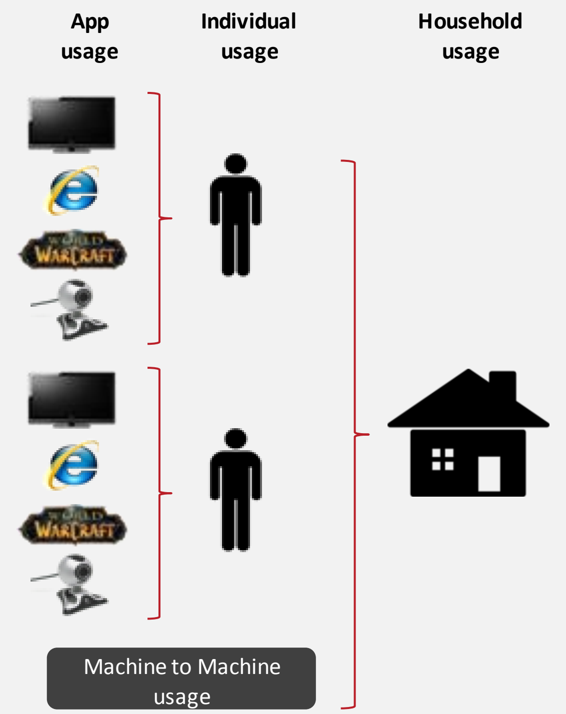

Based on best available data, we have made assumptions regarding both the bandwidth needs and the volume of usage of each of these applications. Note that we have not necessarily based these on today's current usage. In some cases we believe that current usage was constrained by current bandwidth, rather than reflecting what might be reasonably expected absent this constraint. To take one example, while downloading a movie via BitTorrent might typically require a user to wait overnight or longer, we have set expectations to four hours, with this figure decreasing over time so that by 2023 movies can be downloaded in real 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 the 13 most frequent 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[4] and low categories based on their propensity to use the internet. In combination, these splits build to a total of 156 household types.

Figure 2: Household types

13 Demographic types	X	*4 Usage intensities*	X	*3 Main TV types*	=	*156 HH types*
1 adult, 2 adults, 1 adult, 1 child …		High/Med with software DL & video UL/Med/Low		SD/HD/4K		

These 156 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, someone working at home on professional animation might produce very large files for upload, 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.

[4] 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 in the range of 20-30 Mbps or even higher.

Figure 3: 2023 usage profiles, three sample household types

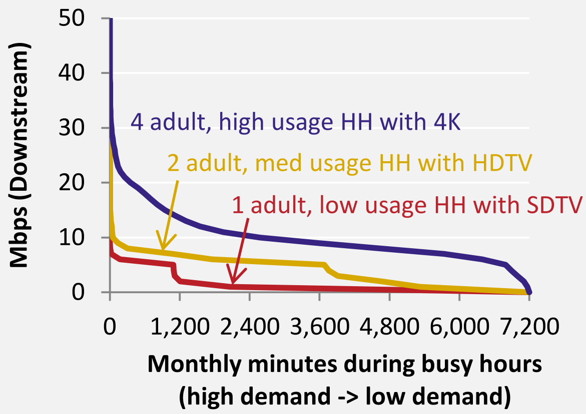

We have used such profiles of household usage to determine '4 minutes excluded monthly' bandwidth demand. By this we mean the bandwidth that would be necessary to serve all but the four busiest minutes in the month (or one per week) – that is, the four minutes at the extreme left of Figure 3. We believe 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.[5]

On this '4 minutes excluded monthly' basis, bandwidth demand grows as shown in Figure 4. As of 2023, the median household requires bandwidth of 19 Mbps (the arrow on the chart), while the top 1% have demand of 35 Mbps.

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

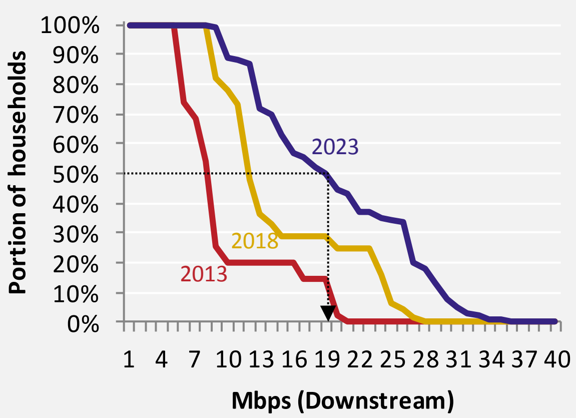

These figures may seem low, particularly by comparison to the results of more informal work in this area. However, the most common type of household comprises just two people. Even if those two are each watching their own HDTV stream, each surfing the web and each having a video call all simultaneously,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

[5] 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.

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. As noted, we have assumed that a Torrent movie download might be in real-time by 2023, 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:[6]

- 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 bandwidth needs, rather than its conclusion. We encourage others to develop their own analyses along the lines set out above, to either validate our work or demonstrate where its conclusions are in error.

[6] 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. A key component of the debate around superfast broadband to date has been about the demand for higher speeds and to what extent this applies to different types of users.

Similarly, an understanding of future needs for bandwidth is important to setting 'threshold' levels – lower bandwidths that might form the basis of a broadband universal service commitment or objective, for instance.[7]

However, while considerable private and public sector investment and attention are being spent on improving broadband, 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 is only used to 0.8% 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.

However, though the *average* utilisation is only 0.8%, this certainly does not mean that 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

[7] The focus of this report is bandwidth, but we note that there are other important technical parameters, such as latency, jitter, packet loss and so on. For some applications these may be more important than raw speed

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.

Figure 5: Busy hours utilisation of UK broadband[8]

Data per connection per month	23 GB
Of which in busy hours (6-12pm)	34%
Data per connection in busy hours	7.8 GB
Traffic per hour in busy hours	43.4 MB
Average usage	0.10 Mbps
Average modem sync speed	12.7 Mbps
Average utilisation	0.9%

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 UK 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 such forecasts, which (to our knowledge) are the first such forecasts in the public domain.

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. 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.

We would welcome any comments you may have, either on assumptions or approach.

8 Communications Chambers analysis of data from Ofcom, *Infrastructure Report 2012 Update*, 16 November 2012. Note that the calculated utilisation also allows for the fact that available bandwidth may be 13-15% lower than sync speed

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.

We thank the Broadband Stakeholder Group for commissioning this piece of work as well as BSkyB, BT, Ofcom, TalkTalk, Three and Vodafone for their additional sponsorship for this project. We are also grateful to the wider group who provided data or comments that improved the inputs of the model. However, the conclusions of this report are Communications Chambers’ own, and do not necessarily represent the views of any of these organisations.

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.[9] 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 the UK has continued to grow, from 27.3m in 2006 to 41.5m in 2013.[11] The portion of households online has grown from 57% to 83% over the same period.

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, stabilising at around 1.9.[12]

Figure 6: Internet users per online household[10]

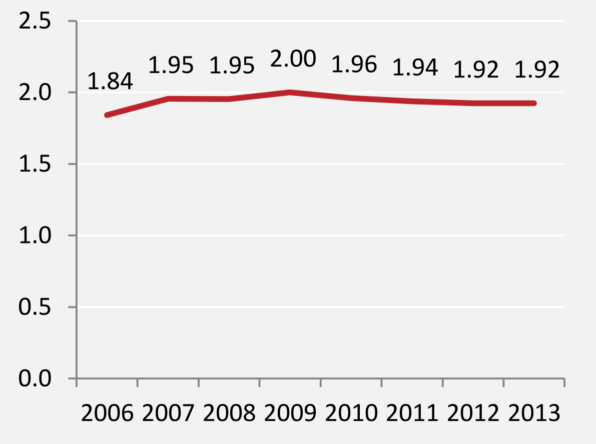

Time spent using the internet

There are obviously natural limits to how much time an individual can spend using the internet. According to Comscore the average online person in the UK currently uses the internet for 37.3 hours per month, but it may be that this figure is approaching saturation. It grew only 5% over the last year. It is only slightly behind North American levels of usage (42.8) and well ahead of the European average (26.9).[13]

---

[9] Though on a standalone basis video may be approaching 'peak bandwidth', given ever improving compression and the limits of the human eye

[10] Department for Communities and Local Government, *Live tables on household projections*

[11] Communications Chambers analysis of ONS, *Internet Access - Households and Individuals*, 2013; individuals who used the internet within the last three months

[12] Ibid.

[13] Comscore, *UK Digital Future in Focus*, 14 February 2013

There is a paucity of data that precisely addresses total time spent online at home across devices - for practical reasons it is difficult to track in-home mobile use. However available evidence suggests that fixed usage growth has been moderate – see Figure 7. Note that the OxIS figures, based on a survey question asking simply about home internet use, may in effect include mobile usage.[15]

Figure 7: Weekly internet usage, hours/user[14]

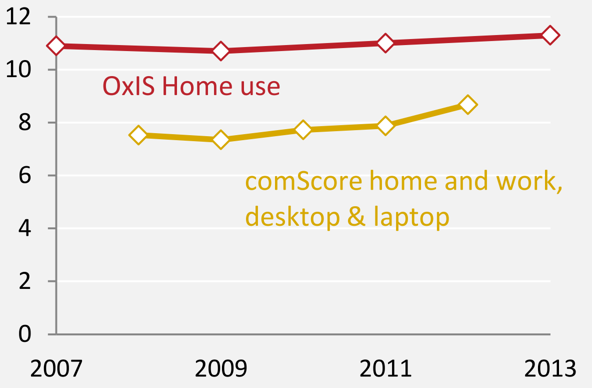

However, mobile devices are undoubtedly providing a boost to usage growth as their adoption increases. Smartphone ownership now stands at 51% and tablet ownership at 24% (a doubling over the past year)[16]. Overall, mobiles and tablets now represent 30.8% of page views, although the figure is likely lower than this overall average for in-home use and higher elsewhere, since out-and-about and at the work place mobile may be the only available option and in many cases a desirable one.

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. Long run figures are not available for the UK. However, they are for a range of other countries (see Figure 8 – note that those figures that are for total traffic are shown with a double arrow, those that are for download only are shown with a downward arrow).

Figure 8: Traffic per fixed BB line, GB/month[17]

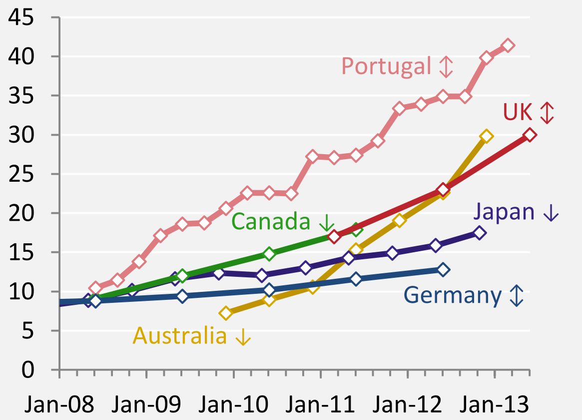

Clearly there is substantial variation between countries. For example, Japan (despite widespread FTTH) has

---

[14] William Dutton & Grant Blank (Oxford Internet Institute), *Cultures of the Internet: The Internet in Britain*, 1 October 2013; Ofcom, *Communications Market Report 2013*, 1 August 2013

[15] That said, we don't think this explains the gap between the two sources, which predates widespread mobile internet. The gap is more likely due to methodological differences

[16] Ofcom, *Communications Market Report 2013*, 1 August 2013

[17] Ofcom, *Infrastructure Report 2012 Update*, 16 November 2012; Ofcom, *Infrastructure Report 2013 Update*, 24 October 2013; BNetzA, *Jahresbericht 2012*, 6 May 2013 ; BNetzA, *Jahresbericht 2011*, 4 May 2012 ; ABS, *Internet Activity, Australia, December 2012*, 9 April 2013 ; ANACOM, *Estatísticas Trimestrais dos Serviços de Acesso à Internet* ; Ministry of Internal Affairs & Communications, *我が国のインターネットにおけるトラヒック総量の把握*, 15 March 2013 ; CRTC, *CRTC Communications Monitoring Report*, September 2012 ; CRTC, *CRTC Communications Monitoring Report*, July 2011 ; CRTC, *CRTC Communications Monitoring Report*, July 2010

relatively low traffic, 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 94 GB for June 2013.[18] OFCA, the Hong Kong regulator, does not break out traffic for residential lines, so this figure may be inflated by business use and as such not a useful comparator.

All countries are seeing growth in per-line traffic, but in general that growth appears linear rather than exponential. The possible exception is Australia, but the acceleration in growth there is likely due to a 90% reduction in usage charges across the market in 2010.[19] (Those charges were previously very high by global standards).

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

[18] OFCA, *Statistics on Customers of Internet Service Providers ("ISPs") in Hong Kong*, August 2013; OFCA, *Customer Access via Broadband Networks*, 27 August 2013

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

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

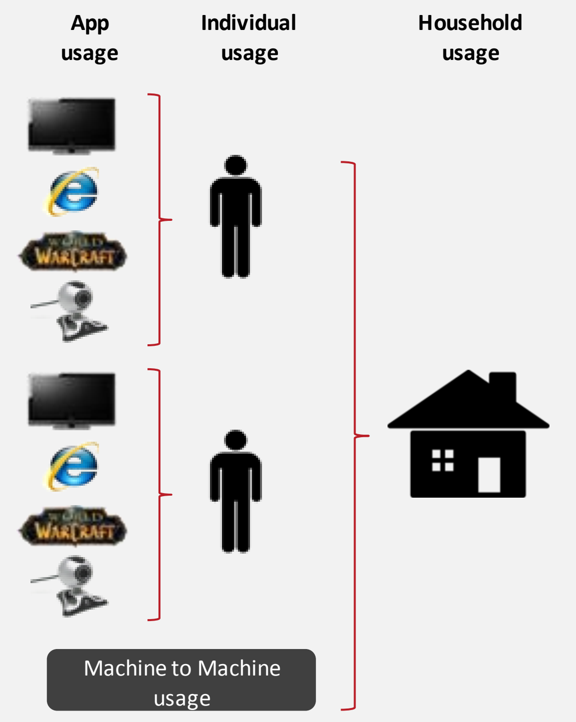

Scope of the model

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 total bandwidth requirement, not how that demand might be met (for instance, via a range of fixed technologies or by wireless).

The model 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 web-usage around the home. Moreover, mobile devices are a key enabler of the combination of applications (such as web usage plus IPTV 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 NHS Direct. 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.

Primary applications are those apps that are primarily used 'one at a time' by a given individual, in particular IPTV 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

Figure 10: Illustrative individual usage over a day

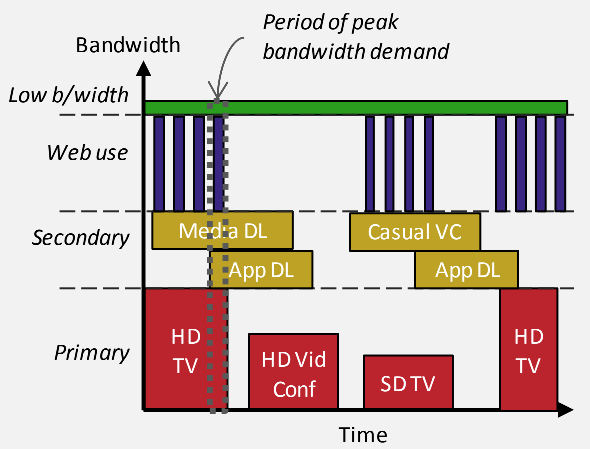

calling that may generally be a primary activity but which, occasionally, might be used in parallel with (say) IPTV consumption.

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
IPTV[20] (SD, HD and 4K)	Cloud storage		[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	P2P / BitTorrent		
Streamed gaming[21] / HD interactive	Mobile OS downloads		
	Software downloads		
	Non-HD video calls		
	Content uploads		

This set of applications includes (but is not limited to) all the key drivers of traffic in the network busy hour today.[22] 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. This compares to the actual network-wide figure of 34% of usage in the

[20] Note that we use IPTV to mean streamed professional content, regardless of whether delivered to a TV set or a mobile device

[21] 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

[22] Based on Sandvine, *Global Internet Phenomena Snapshot: 1H 2013 - Europe, Fixed Access*,13 May 2013

six busiest hours.[23] Our figures imply much greater concentration, but that is appropriate since traffic is likely to be more concentrated at the household level than it is at the network level.[24]

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 IPTV.

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.

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. 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

[23] Ofcom, *Infrastructure Report 2012 Update*, 16 November 2012

[24] 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

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 IPTV streams will tend to overlap, since user A watching IPTV may encourage user B (with different tastes) to seek out their own IPTV 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 IPTV stream, it may be less likely that person B's own IPTV usage is simultaneous, since they may instead be watching the first stream with person A. (Or person A may be watching IPTV precisely because person B is watching linear TV –without using bandwidth - on the main set. Linear is still by an enormous margin the predominant mode of TV consumption).

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 IPTV 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 IPTV 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.

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 – 13 types of household); whether they use SD, HD or 4K TV (3 types) and whether the individuals in the household are light, medium[25] or heavy users of the internet (4 types). Since these categories multiply up, the model considers 156 different household types overall, enabling us to understand not just what

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

the typical household requires, but also to understand the spread from the most demanding to the least demanding household type.

Demographic types

Based on ONS statistics,[26] we have estimated the mix of household types as shown in Figure 12. Note the importance of one and two person households which represent 64% of the total.

Figure 12: Modelled household demographics

		Children			
		0	1	2	3+
Adults	1	29.0%	3.2%	1.8%	0.6%
	2	31.8%	6.5%	8.9%	2.7%
	3	8.2%	1.2%	2.6%	1.0%
	4	2.7%			

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 have increased across all demographics. For similar reasons, we 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 aged 0-15, and thus this category contains a significant number of the very young, who will pull average usage down. In this age band, 27% are aged three or less.[27]
- 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 PC time for those aged 2-15 was 30% of those aged 16+ in August 2013[28] (although we would expect children to be heavier users of tablets and mobiles in the home offsetting this somewhat).

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 21% of the population, and are only present in 28% 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

[26] ONS, *Families and Households 2012*, 1 November 2012
[27] ONS, *2011 Census - Population and Household Estimates for England and Wales, March 2011*, 16 July 2012
[28] Communications Chambers analysis of Nielsen Netview data, August 2013 (home only)

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 IPTV can create a high bandwidth need, but the take-up of such devices across most households is uncertain.[29]

Development over time

The model develops a ten year forecast for unconstrained[30] 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 IPTV)
- 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.

[29] 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

[30] 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.[32]

Figure 13: Video & audio bandwidth per Canadian Radio-television and Telecommunications Commission[31]

Category	Application	Average speed (Mbps)	
		Down	Up
Video streaming	Netflix (default)	0.657	0.084
	Netflix (good)	0.691	0.083
	Netflix (better)	1.343	0.162
	Netflix (best)	4.866	0.512
	YouTube (HD 720p)	1.537	0.173
	YouTube (HD)	2.522	0.298
	YouTube (SD)	0.443	0.063
Audio streaming	Grooveshark	0.224	0.034
	Slacker	0.132	0.035
Real-time	Skype (w. video)	0.237	0.237
	Skype (audio only)	0.042	0.042
	Google Talk (w. video)	0.263	0.263

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

[31] CRTC, _Communications Monitoring Report 2012_, 5 September 2012
[32] 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 2/3 of consumer traffic in the UK (and rising) [34].

In this primary category the model considers: IPTV (in various resolutions), YouTube and similar video; HD video calls; and streamed gaming.

Figure 14: Total traffic by type in petabytes (PB), 2013 (UK consumer applications)[33]

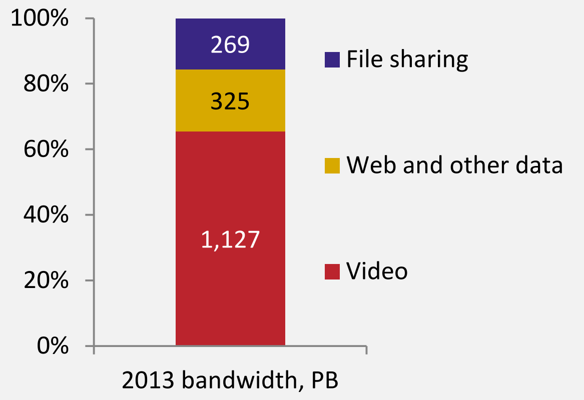

IPTV viewing

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

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

To estimate the total amount of IPTV, we built up from current usage of iPlayer.

BBC iPlayer received 181m requests for TV programmes in May 2013.[35] At an average stream play time of 22 minutes[36] this equates to total monthly iPlayer viewing of 66.37m hours with a monthly active online audience of 43.6m.[37] This implies that the average online individual watches 91 minutes of iPlayer a month (including some non-home usage).

Given that Netflix, LoveFilm and 4oD combined account for approximately the same traffic in the UK as iPlayer,[38] the total amount of viewing of these services is in the region of 182 minutes. To allow for a long tail of other services, we have 'rounded up' this figure.

We assume that in 2013, the average IPTV viewing is 210 minutes per use per month.

While we believe that live broadcast TV will remain the dominant mode of viewing for the foreseeable future, the amount of IPTV is

---

[33] Cisco, *Visual Networking Index*, accessed August 2013 (UK 2013 estimate). Note that 3 PB of gaming traffic has been excluded from the chart

[34] Cisco, *Visual Networking Index*, accessed August 2013

[35] BBC, *iPlayer Monthly Performance Pack*, May 2013

[36] BBC, *BBC iPlayer Behind The Scenes*, 2009

[37] Ofcom, *Communications Market Report 2013*, 1 August 2013

[38] LSE, *European Internet Traffic: Indicator of Growth and Competition in Digital Services – A summary*, 21 Feb 2013 [referencing Sandvine data]

likely to grow considerably from its current 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.

One early indicator of the potential is Netflix. It had 1.5-2m subscribers in the UK in August 2013, [39] despite only launching in early 2012. According to Netflix's 2012 annual report, they had "33 million members in over 40 countries enjoying more than one billion hours of TV shows and movies per month".[40] This equates to approximately 30 hours per subscriber per month, or around an hour a day. Some of this viewing is likely to be shared, [41] and early Netflix adopters are likely to be higher than average users of IPTV. Nevertheless it illustrates the possibility of material growth in IPTV viewing.

We assume that IPTV viewing will grow by an average CAGR of 25%.

This growth rate is roughly equivalent to iPlayer's recent average annual growth in requests, though it is perhaps conservative to assume that this growth could sustain over the long term. It is also worth noting how important the adoption of mobile devices - and in particular tablets - have been to iPlayer's recent growth. (See Figure 15 – note that there is strong seasonality in iPlayer consumption, which typically falls off in the summer months).[43] It may be therefore that IPTV growth will slow as tablet adoption reaches saturation, though that point is likely some years off.

Figure 15: iPlayer TV requests (m) by platform[42]

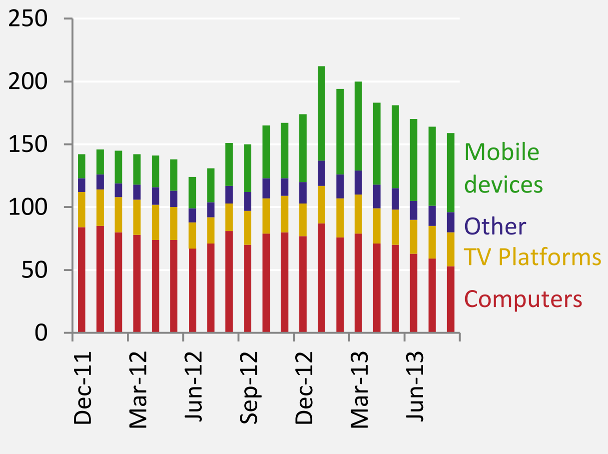

Our 25% growth assumption implies that IPTV will be 1,956 minutes per person per month by 2023 (64 minutes per day). IPTV in 2013 would therefore represent 27% of total TV viewing based on today's viewing habits.[44]

---

[39] Enders Analysis report referenced by The Guardian, *Netflix reaches 1.5m UK subscribers for its internet video service*, 21 August 2013

[40] Netflix, *Annual Report 2012*

[41] If on average 1.5 people watched a Netflix stream, this corresponds to 40 minutes of Netflix viewing per person per day

[42] BBC, *iPlayer Performance Packs*. Note that a measurement problem was fixed in February 2012 which resulted in an increase in reported mobile usage from that month onwards

[43] This is roughly consistent with the seasonality of the UK internet as a whole – broadly speaking, traffic levels are flat during H1 of each year, with all growth occurring in H2 after the summer. See for instance the statistics for traffic through the London Internet Exchange: LINX, *LINX Aggregated Traffic Statistics* [accessed 14 October 2013]

[44] Average TV viewing is 241 minutes per day, according to Ofcom, *Communications Market Report 2013*, 1 August 2013

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.[45] 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 IPTV viewing

A significant driver of IPTV bandwidth requirements will be the resolution and format of the video. Main set viewing is currently migrating from standard definition to HD, and may in time shift to 4K.

The proportion of HD households in the UK (households with a HD-capable TV set and access to HD channels) grew to 42% in 2012[46] with the trend likely to continue.

> We have assumed that the proportion of HD households is 50% in 2013, growing at 4% 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. 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"[47].

> We have assumed that HD viewing represents 20% of IPTV 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 8% in 2020.[48]

> Based on Futuresource's forecast, we have assumed that 4K penetration will grow from 0% in 2013 to 19% in 2023.

---

[45] See for example: King's College London, *'Intelligent and green' iPlayer records favourite TV in advance and reduces internet traffic*, 17 June 2013
[46] Ofcom, *Communications Market Report 2012*, 18 July 2012
[47] Farid Husseini, *High Definition Television: It's a must have, but is it also a must see?,* 11 July 2012
[48] Futuresource, *Ultra HD: The Content Perspective,* 27 June 2013

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 IPTV viewing in 4K households by 2023.

SD, HD and 4K IPTV 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[49], 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 2.2 Mbps and 2.6 Mbps in July 2013[50]. This average 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).
- 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)[51].

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.

- YouView recommends a minimum broadband speed of 3 Mbps for a service which can stream HD content[52].

[49] Elena Puigrefagut for the EBU, *HDTV and Beyond*, 2012
[50] Netflix, *UK ISP Speed Index*, August 2013
[51] BBC, *iPlayer FAQs*, accessed August 2013
[52] YouView, *Broadband speed FAQs*, accessed August 2013

- BBC iPlayer requires 3 Mbps for HD content[53] based on 720p25 (720p resolution as 25 Hz)
- According to Netflix, 'Super HD' (in 1080p standard) requires between 5 Mbps and 7 Mbps[54]
- The EBU estimate that the bitrate for an HD programme varies between 7.85 and 10.85 Mbps[55].

We have assumed the HD video requires 5 Mbps, and 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:

- Deloitte state that by 2014 4K TV will require 20 Mbps[56]
- *Satellite Executive Briefing* reports that the implementation of HEVC[57] will reduce requirements from 48 Mbps to "24 Mbps at the most"[58]
- Eutelsat suggest less than 20 Mbps for 4K at 50 fps[59] using MPEG-4, falling to 12 to Mbps with HEVC
- Netflix CEO Reed Hastings suggests "around 15 Mbps … If you've got a 50-megabit connection you'll be fine"[60]
- Encoding technology supplier ATEME state: "We're showing 4K at 60 Hz with excellent quality between 11 and 15 megabits per second"[61]
- 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"[62]

In order not to underestimate demand and to allow for particularly demanding content such as live sports, we have assumed that 4KTV requires 30 Mbps in 2013.

---

[53] BBC, *iPlayer FAQs*, accessed August 2013
[54] Netflix, *Netflix Super HD*, accessed September 2013
[55] Elena Puigrefagut for the EBU, *HDTV and Beyond*, 2012
[56] Deloitte, *Survival of the fastest: TV's evolution in a connected world*, September 2013
[57] High efficiency video coding, a recent video compression standard now appearing in consumer equipment
[58] Elisabeth Tweedle, "4K TV: A technology push or a demand pull?", *Satellite Executive Briefing*, September 2013
[59] Frames per second
[60] DSL Reports, *Netflix: 4K Video Will Need At Least 15 Mbps*, 23 September 2013
[61] Fred Dawson, "Underestimating 4K Potential Poses Mounting Risks for SPs", Screen Plays, 7 October 2013
[62] "HEVC goes beyond HD", TVBEurope, 4 June 2013

In our modelling approach we have not explicitly addressed 3DTV 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 IPTV are likely to remain minimal. CRTC[63] found that even the highest quality Netflix and YouTube video streams had an average upstream bandwidth of 30 – 50 kbps.

We have assumed a slightly higher 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).

Estimates for the rate of improvement in compression vary, although tend to be around 10%:[65]

Figure 16: Required video bandwidth over time[64]

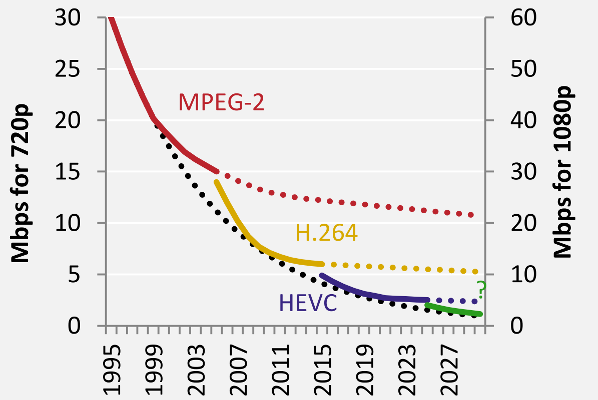

- Sky suggest that the bitrate requirement for HD content has fallen from around 30 Mbps in 2002 to 10 Mbps in 2010[66], equating to an annual decline of around 13%.
- Zetacast suggest that video bandwidth requirements halve every seven years (see Figure 16), corresponding to an average annual decline of around 9%.
- According to Motorola, encoder technology evolution has resulted in a 10-20% annual improvement in efficiency[67].
- A 2006 University of Essex study predicted annual reduction in bitrates of 7%[68] (although with improving quality).

---

[63] CRTC, Communications Monitoring Report 2012, 5 September 2012
[64] ZetaCast, Technical Evolution of the DTT Platform, 2012
[65] 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
[66] Sky, Beyond HD Masters 2013, 2013
[67] Motorola, Opportunity and impact of video on LTE Networks, 13 May 2013
[68] Ghanbari et el, Future performance of video codecs, 13 July 2006

We have assumed an annual improvement in compression of 9% for SD, HD and 4K TV.

YouTube and other non-IPTV 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 24% of peak downstream traffic[69]. Since the characteristics of its traffic are materially different to that of IPTV, we have considered it separately.

Estimates of YouTube use tend to be around 5-6 hours per user per month:

- In February 2013, UK internet users spent 203m hours watching YouTube content[70] which, based on Netview's estimate of 40.2m active users that month[71], corresponds to 303 minutes per user per month
- In the US, the average YouTube user spent 311 minutes watching YouTube content in May 2011[72]
- According to YouTube, their global audience of 1bn users watch 6bn hours of content each month (360 minutes per person per month)[73]

These estimates will include out-of-home use and therefore over-estimate the YouTube traffic passing through the home network.

We have assumed the average adult watches 300 minutes of non-IPTV video per month from the home, growing at 13% per year

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

- 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[74].
- Ameigeiras et al's analysis of a sample of YouTube videos[75] found overall average bitrates of 510 kbps, and averages by resolution and format of 301 kbps (240p), 527 kbps (360p)

---

[69] Sandvine, *Global Internet Phenomena Snapshot: 1H 2013 - Europe, Fixed Access*,13 May 2013
[70] Experian, *UK Internet users make over a billion visits a month to video sites*, 26 March 2013
[71] Nielsen, Communications Chambers analysis of Netview data, accessed September 2013
[72] Techcrunch, *comScore: The Average YouTube Viewer Watches 5 Hours Of Videos A Month*, 17 June 2011
[73] YouTube, Statistics, accessed September 2013
[74] AdTerras Per Aspera blog, *Approximate YouTube bitrates*, 24 May 2010
[75] Ameigeiras et al, *Analysis and modeling of YouTube traffic*, 26 June 2012

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 IPTV 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 includes HD video for the iPhone 5 and fourth-generation iPad[76], for example. Some smart TVs also come with Skype capability.

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.[77] Apple iChat requires up to 900 kbps for 'best' quality.[78]

Three, four and more way sessions may also increase the speed requirements further. For example, Apple state a 4-way video conference at 'best' quality requires 1.8 Mbps for the initiator, and 300 kbps for 3 participants.[79]

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 IPTV standard. The CRTC found that Skype and Google Talk video had an average speed of 200 –

---

[76] PC Mag, *Skype Update Adds HD Video Calling to iPhone 5, 4th-Gen iPad*, August 12 2013
[77] Skype, *How much bandwidth does Skype need?,* accessed September 2013
[78] Apple, *iChat system requirements*, accessed September 2013
[79] Apple, ibid

300 kbps, considerably lower than SD video requirements of 1.5 – 2.5 Mbps. [80]

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,[81] 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 – an NIH study achieved good results using 3G macrocellular networks. [82]

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 last year 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, 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 320 minutes per user per month by 2023.

[80] CRTC, *Communications Monitoring Report 2012*, 5 September 2012
[81] WiK, *Market potential for high-speed broadband connections in Germany in the year 2025*, 15 January 2013
[82] Locatis et al, *Video Medical Interpretation over 3G Cellular Networks: A Feasibility Study*, December 2011

In terms of bandwidth requirements, Microsoft recommended a minimum connection speed of 3 Mbps for Xbox Live gaming[83] (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.[84]

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 175m users globally, and 1bn files saved each day[85].

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 download speed, and uploads are automatically throttled to 75% of the maximum upload speed to prevent slowdown in browsing[86]. 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

[83] Microsoft, *Xbox Live support*, accessed September 2013
[84] Onlive *Tech Specs*; Time, *Sony Bets On Cloud Gaming with Gaikai Purchase, but Don't Expect Drastic Changes*, 3 July 2012
[85] The Guardian, *Dropbox passes 175m users and turns up heat on Apple iCloud with new APIs*, 10 July 2013
[86] Dropbox, *FAQs*, accessed September 2013

et al[87] (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), assume a rise in penetration of cloud storage services from 20% today to 75% in 2023, and an increase in the average number of shared devices per user (from 1.2 today to 2.0 in 2023). [88]

We note that some claim extremely high bandwidth requirements for cloud computing. WIK suggests speeds of up to 100 Mpbs may be required. [89] 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". [90]

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. 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 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 of 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.

---

[87] Drago et al, *Inside Dropbox: Understanding Personal Cloud Storage Services*, 14 November 2012
[88] Drago et al, *Inside Dropbox: Understanding Personal Cloud Storage Services*, 14 November 2012
[89] WiK, *Market potential for high-speed broadband connections in Germany in the year 2025*, 15 January 2013
[90] FTTH Council, *Q&A Webinar on WiK report*, 15 January 2013

Content downloads

As well as streaming video, consumers may also download content, such as TV programmes, films and music files.

To estimate the frequency and scale of content downloads, we use Kantar Media's report for Ofcom on the number of legal content downloads in the UK through services such as Apple's iTunes store. [92] (We use Kantar's estimate of illegal downloads as a basis for estimating the scale of torrenting, discussed later.)

Figure 17: UK consumer-reported legal content consumption[91]

	Total (m/month)	Per downloader (/month)
Music tracks	340.0	8.5
TV programmes	91.7	2.3
Films	24.3	0.6

Kantar's analysis suggests that 8.5 songs, 0.6 films and 2.3 TV programmes were legally downloaded per online user per month. Assuming a film is 2 hours, and a TV programme 45 minutes, this is around 187 minutes of content per user per month.

We have assumed that the average user downloads 180 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 around 11% of internet users legally download content each month today,[93] we assume this rises to 50% penetration by 2023.

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.

[91] Adapted from Kantar Media, for Ofcom, *Online Copyright Infringement Tracker Wave 3 (Covering period Nov 12 – Jan 13) - Overview and key findings*, 28 May 2013. Note that these figures are for legal consumption which we assume is analogous to content downloading
[92] Ibid.
[93] Ibid.

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. It is relatively popular, with 4m UK users each month in 2012,[94] and it is used for large files. Europe wide, it represents 12.2% of peak period downstream traffic.[95] However, this is a percentage that has dropped in recent years (though may have stabilised).

While some have claimed this drop is due to wider availability of legal downloads,[96] it is likely largely due to 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 do 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].[97]

[94] Nielsen, quoted in *BPI, Digital Music Nation 2013*, January 2013
[95] Sandvine, *Global Internet Phenomena Snapshot: 1H 2013 - Europe, Fixed Access*,13 May 2013
[96] Netflix have claimed that "when we launch in a territory the Bittorrent traffic drops as the Netflix traffic grows". Stuff, *Netflix's Ted Sarandos talks Arrested Development, 4K and reviving old shows*, 1 May 2013
[97] About.com, *How to download with BitTorrents*, August 2013.

By their nature, P2P downloads are constrained by the 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.

Figure 18: UK consumer-reported illegal content consumption [98]

	Total (m/month)	Per downloader (/month)
Music tracks	93.3	17.0
TV programmes	17.3	3.2
Films	9.7	1.8
e-books	6.0	1.1
Computer software	2.3	0.4
Video games	2.3	0.4

In considering volumes of usage, we have again drawn on Ofcom's Copyright Infringement Tracker. This is *not* because BitTorrent and piracy are synonymous – the protocol has many legitimate users (for example, distribution of Linux) and the company BitTorrent Inc is itself legitimate. However, we believe that piracy is still the key volume driver for use of the protocol.

Ofcom found that 18% of those online how downloaded illegal content. Most important in terms of bandwidth is likely the 3.2 TV programmes and the 1.8 films each such downloader consumed (at 45 minutes a TV programme and 120 minutes a film, this equals to 360 minutes of content).

We have assumed that the average user uses P2P to download 360 minutes of content per month in 2013, with a 10% CAGR.

Today, it is the SD versions of such files that are more popular, with (according to one study) 92% of users choosing them.[99] 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.

In determining bandwidth required, we have assumed that:

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

[98] Adapted from Kantar Media, for Ofcom, *Online Copyright Infringement Tracker Wave 3 (Covering period Nov 12 – Jan 13) - Overview and key findings*, 28 May 2013. Note that these figures are for illegal consumption (including streaming), but we are assuming streamed illegal content is a rarity

[99] Petrus Potgeiter, *High-definition content and file sharing networks*, 18 November 2012

Non-HD video calls

As well as HD video calls, which are assumed to be a primary application, consumers may also multi-task video calls with other activities (for example chatting on Skype while watching IPTV).

Skype currently has 280m users worldwide,[100] generating 2bn minutes of traffic per day,[101] or 214 minutes per user per month. Based on an estimated penetration of Skype (and similar services) of 25% we have therefore assumed:

> The average individual has 54 minutes of non-HD video calls in 2013, growing to 5 hours by 2013.

Note that we believe the assumption of 5 hours is conservative, even allowing for the possibility of video windows being left open for telepresence. Five hours per month is equivalent to combined fixed *and* mobile voice minutes per capita today.[102]

For one-on-one video calls, a high quality Skype video call recommends 500 kbps[103] downstream and upstream capacity. As with other forms of video we have treated in the model, we assume this falls due to 9% annual compression improvements.

Software downloads

Software downloads are particularly challenging to address in the context of required bandwidth. The files involved can be large – OS X Mountain Lion is 4.5 GB, and console games are typically 6-9 GB.[104] We will use Apple OSs and console games as our case studies, since we believe they represent are 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.[105]

---

[100] Skype, *Skype and Messenger Coming Together: The Next Chapter*, 15 February 2013
[101] Skype, *Thanks for Making Skype a Part of Your Daily Lives – 2 Billion Minutes a Day!*, 3 April 2013
[102] Communications Chambers analysis of data in Ofcom, *Telecommunications Market Data Update Q1 2013*, 22 August 2013
[103] Skype, *How much bandwidth does Skype need?*, accessed September 2013
[104] Deloitte, *Media Consumer Survey 2013*, 30 April 2013.
[105] Gigaom, *Lion download too fat? There's Apple store (WiFi) for that*, 21 June 2011

However, downloads are relatively infrequent. We estimate the average household downloads 0.6 console games per year.[106] Apple updates its computer OS roughly annually, and has approximately an 11% share in the UK, implying 0.11 downloads per year per household (assuming 1 Mac per home).[107] For whatever reason, new Windows OSs appear to be much less likely to be downloaded.

One reason for infrequent games downloads is that there is still a strong preference for physical copies - 81% prefer physical over digital copies, and this implies only 10% of households (approximately) download today.[108] 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.[109]

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 Mountain Lion, some users reported downloads in the range of 1 to 3 hours,[110] others as quick as 30 minutes.[111]

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. [112] The new Xbox will also allow 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".[113] (The PlayStation 4, also coming out later this year, will have the same capability.)[114]

Furthermore updates may take place outside of peak hours, as it already does with Windows. Microsoft's Windows Update client

---

[106] Communications Chambers analysis of figures from Deloitte, ibid
[107] *Statcounter*
[108] Deloitte, *Media Consumer Survey 2013*, 30 April 2013.
[109] Deloitte, ibid, and Hybris, *PC and Console Gaming: The end of the DVD?*, 9 October 2012
[110] Tuaw, *How long will it take to install Mountain Lion?*, 25 July 2012
[111] Mashable, *Apple OS X Mountain Lion: The Good, the Awesome, the Could-Do-Better*, 25 July 2012
[112] IGN, *Ask Microsoft anything about Xbox One*, 5 August 2013
[113] Engadget, *Xbox One has play as you download functionality similar to PlayStation 4*, 16 July 2013
[114] Engadget, *PlayStation 4's UI and inner workings detailed: No more booting, games download as you play them*, 20 February 2013

"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."[115]

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:

- Five downloads per year of 7GB (8% CAGR), but initially only in the 20% of households that are 'high use'. 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

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.[116] (Google's Android Jelly Bean update for the Samsung S4 smartphone was considerably smaller at 151MB[117]).

We have assumed a mobile OS update download size of 700MB, growing at 10% each year. We further assume that such updates are annual.

OS updates will apply to all those with smartphones (around 60%[118] 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 a key driver for bandwidth requirements.

---

[115] Microsoft, *Windows Update Explained*, September 2008
[116] The Telegraph, *iOS 7: how to install the new software for iPhone and iPad*, 19 September 2013
[117] Technostop, *Samsung galaxy S4 Jelly Bean 4.3 Update*, 3 August 2013
[118] Ofcom, *Communications Market Report 2013* stated 51% in Q1 2013

We have assumed a 30 minute download time in 2013, falling to 15 minutes in 2023 (for an update that is almost 2GB)

We note that current download times may be much longer than 30 minutes, with the recent iOS download taking "many hours" for some[119]. Particularly immediately after release, server capacity appears to be a constraint for iOS downloads.

Content Upload

By some margin, the most popular site for photo sharing is Facebook.[120] 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).[121] The average size of each photo is just 107 KB, likely because the Facebook app compresses photos before uploading them.[122] 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 yet a widespread activity. In the UK only 22% of internet users have ever uploaded a video.[123] YouTube's reach in the UK is 40%, though of course this includes individuals who only view.[124] Globally, YouTube receives 100 hours of video every minute and has 1 billion visitors per month, 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).[125]

For modelling purposes, we assume that initially 20% of UK internet users are uploaders, at an average of 10 minutes per month. Over time, the number of users increases to 40%

This volume of video uploads 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.

---

[119] The Guardian, *iOS 7 download delays frustrate Apple customers*, 19 September 2013
[120] Mary Meeker, *Internet Trends*, 29 May 2013
[121] Facebook, *Form 10-K*, 1 February 2013
[122] Communications Chambers analysis of figures from Dhruba Borthakur, *Petabyte Scale Data at Facebook*, November 2012
[123] Ofcom, *Adults media use and attitudes report 2013*, April 2013
[124] UKOM
[125] YouTube, *Statistics*

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.[126] 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.[127]

- 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 UKOM, the average user spent 2,238 minutes per month online through a PC[128]. Assuming a 45% uplift to reflect non-PC use[129], this corresponds to 3,243 minutes of web use per user per month. This will double count some time spent on primary and secondary apps, in particular short-form video such as YouTube, although (conservatively) we have made no adjustment for this.

Data from UKOM, Netview and Ofcom suggests that annual growth in per-person web use has been around 5%. However as general web use continues to mature, long-run growth will be slightly lower. 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.

[126] GSMArena, *Apple iPhone 5 vs. Samsung Galaxy S III: All rise*, 25 October 2012
[127] YouTube, *Video stuck in processing*
[128] Communications Chambers analysis of UKOM data
[129] Based on data in Ofcom, *Communications Market Report 2013*, 1 August 2013

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

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.

Figure 19: Web usage and page loads (illustrative)

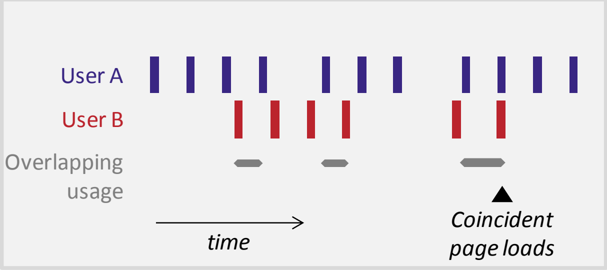

In practice, the experience of web-surfing is driven at least as much by latency (the round-trip time for a request for data and the response) as it is by bandwidth (how much data can pass simultaneously through a particular element of the network).[130] Indeed, beyond a certain point increased bandwidth makes relatively little difference. For instance, according to one study, increasing bandwidth from 1.5 Mbps to 5 Mbps might only reduce page load times by only 12%.[131] Others have found diminishing returns from 2 Mbps onwards.[132]

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

We have also assumed:

- Total page weight rises to from 4.0 MB to 6.4 MB in 2023 (which increases how long the provided bandwidth must be sustained)
- Average time spent on a page is 53 seconds[133]
- The downstream to upstream ratio for web traffic is 7.58[134]

Note that the above assumptions (in particular page weights) are based on the web as accessed from fixed devices. We believe this approach is highly conservative. Increasingly the web is accessed from mobile devices, with 24.0% of UK page views coming from

[130] Note that different access technologies have different latencies – generally mobile will have the highest and FTTH the lowest
[131] NCC Group, *Will faster user bandwidth fix your website performance woes?* 5 June 2013
[132] Ilya Grigorik, *Latency: The New Web Performance Bottleneck*, 19 July 2012
[133] Communications Chambers analysis of Nielsen Netview data, July 2013
[134] Sandvine, *Global Internet Phenomena Snapshot: 1H 2013 - Europe, Fixed Access*,13 May 2013

phones[135]. 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 screen size and mobile data caps. This may greatly reduce effective average page weights for surfing, but this effect is not factored into our forecast.

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 a downstream bandwidth of 1 Mbps, and 1% occupancy during busy hours (growing to 3% by 2023).

Upstream requirements will also be trivial. 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[136], *World of Warcraft* 2 kbps and even high intensity *Halo 3* requires 60 kbps.[137]

[135] Chris Thomson (Comscore), *Europe Digital Future in Focus 2013*, March 2013
[136] Saldana and Suznjevic, *Tutorial: Traffic of Online Games*, 1 August 2013
[137] Andreas Petlund et al, *Network Traffic from Anarchy Online: Analysis, Statistics and Applications*, 22 February 2012

Voice calls

These require 100 kbps or less.[138]

Smart metering

Bandwidths are often in the tens of Kbps or less. Italy's 30m smart meters use a bandwidth of 2.4 kbps.[139]

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"[140].

Remote working via VPNs

Citrix estimate that their popular XenDesktop virtual desktop / remote working service requires an average bitrate of just 78 kbps[141]. 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.

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 20 – note the log scale on the horizontal access. The highest bandwidth application, software downloads at 21 Mbps, has relatively brief demand, at 20 minutes per month. 4K TV, at 12 Mbps, has more sustained demand, at 100 minutes per month.

Figure 20: 2023 usage and bandwidth by app

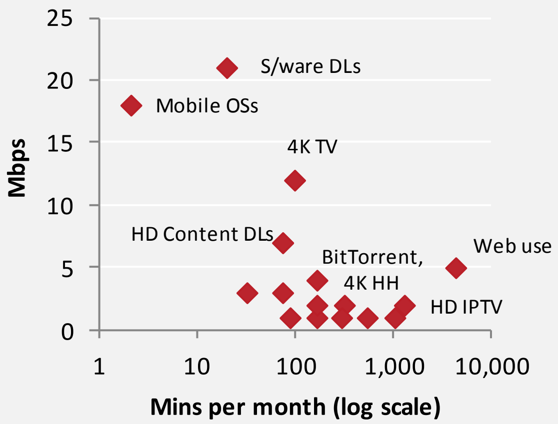

The heaviest application - as measured by time - is web usage, with 4,360 minutes per month per adult (or 2.42 hours per day). While not the highest, its bandwidth requirement is still appreciable at 5 Mbps.

The applications that are most significant for peak demand are, broadly, those to the right or to the top of Figure 20. Heavily used apps such as surfing and IPTV create a 'base load' of bandwidth

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

[139] 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

[140] Harrop and Armitage, Quantifying the Broadband Access Bandwidth Demands of Typical Home Users, December 2006

[141] Citrix, XenDesktop Planning Guide: User Bandwidth Requirements

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 in the month that set peak bandwidth demand.

We now turn our attention to bandwidth demand results implied by these 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 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 annual figures, 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) 50 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[142] 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.[143] 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.[144] 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

[142] 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
[143] Roughly equivalent to excluding one peak minute per week
[144] 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 'four nines' 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.

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 156 different household types:

Figure 21: Household types

13 Demographic types	X	*4 Usage intensities*	X	*3 Main TV types*	=	*156 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

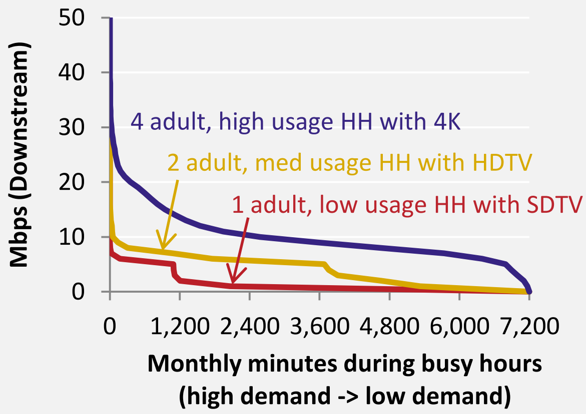

In the 1 adult, low usage household, the broadband connection is idle for much of the month, requiring 0 Mbps. At the other extreme, the 4 adult, high usage household with a 4K TV has demand almost constantly during the busy hours, frequently above 10 Mbps and occasionally above 30 Mbps.

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 websurfing (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 38 Mbps, though note that in the excluded peak minutes, the bandwidth requirement reaches over 50 Mbps.

Figure 23: 2023 usage profiles, three sample household types

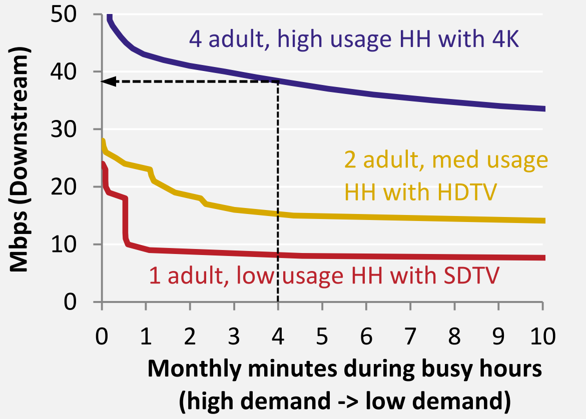

There are three important points to note regarding this approach. Firstly, while we have ranked minutes and considered the small number of busiest minutes, these minutes need not be continuous – indeed, they are likely to be widely scattered throughout the month. Further, given that much traffic (particularly websurfing) 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 10 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 156 household types, we are in a position to aggregate these to understand the overall spread of demand.

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

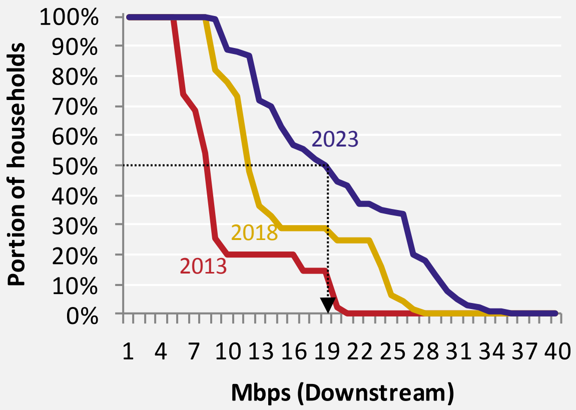

Figure 24 shows the distribution of demand over time. The curves show the percentage of households that have at least a given level of demand (on a '4 minutes excluded' basis). The arrow indicates that in 2023, 50% of households have a demand of 19 Mbps or more. (Put another way, the median demand is 19 Mbps). In that year 10% have a demand of 30 Mbps or more, and the top 1% have a '4 minutes excluded' demand of 35 Mbps.

As would be expected, these are higher levels of demand than in earlier years – the median demand in 2013 and 2018 is 8 and 12 Mbps respectively.

We note that an increase in median demand of 11 Mbps (from 8 to 19 Mbps) over the next ten years compares to the growth *to* 8 Mbps over the previous twenty years.[145]

A 2023 median demand of 19 Mbps may seem low, but needs to be seen in the context of the continuing benefits of video compression, and the fact that 64% 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 15 Mbps in 2023.[146]

Figure 25 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).

[145] 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)

[146] 2 x HD at 2 Mbps, 2 x surfing at 5 Mbps and 2x video call at 200 kbps

Figure 25: 2023 Downstream demand by minute and household

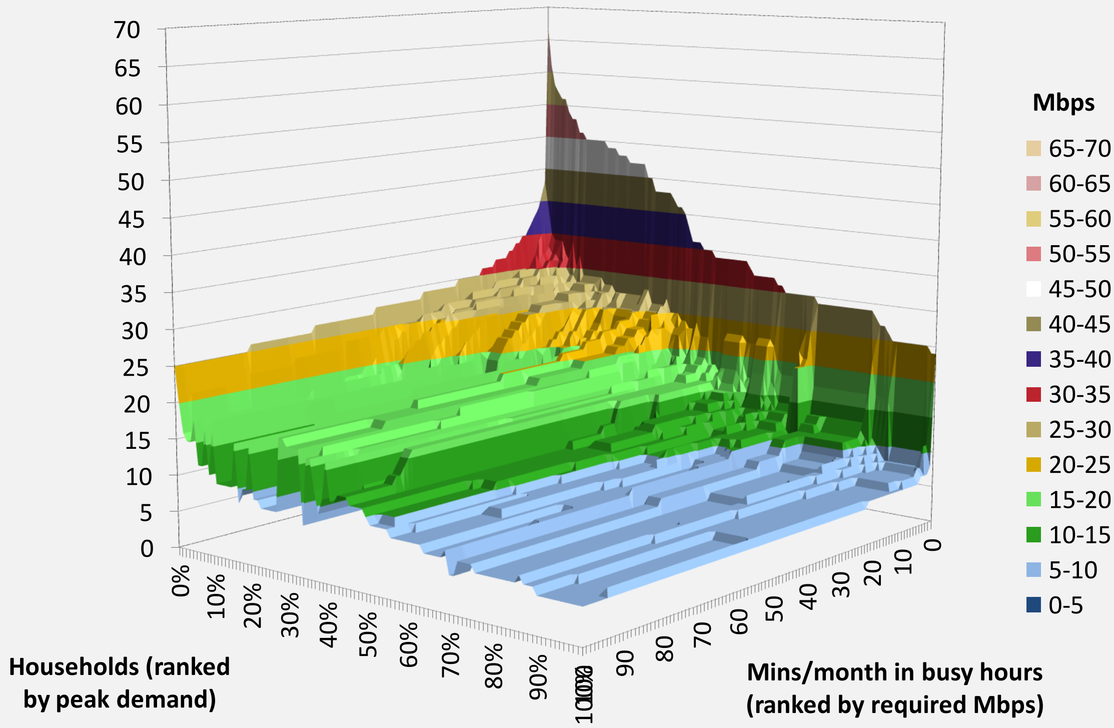

For all households there is the possibility of high demand (the 'cliff face' at the back right). However, the expected duration of that demand is short, resulting in the sharp drop off.

In the back corner, the most demanding households trigger some demand above 50 Mbps (and even 60 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, IPTV creates much downstream demand, but drives little upstream traffic. Uploading video (to YouTube, for example) is the reverse.

Median upstream demand grows from 1.1 to 2.4 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.

Figure 26: 2023 Upstream Bandwidth demand distribution (4 mins excluded basis)

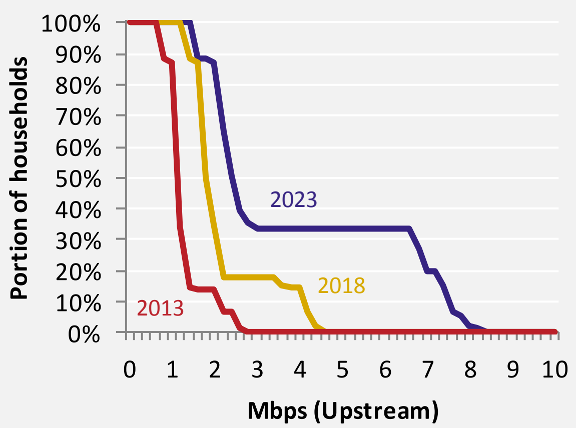

Sensitivities to model results

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 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 27 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 19 to 24 Mbps, since the additional 3 minutes brought into scope are very busy. Conversely, if the criterion is loosened to 30 minutes (equivalent to one degraded minute per household per day), the median required bandwidth drops to 12 Mbps.

Figure 27: 2023 Downstream Bandwidth demand distribution (various mins excluded basis)

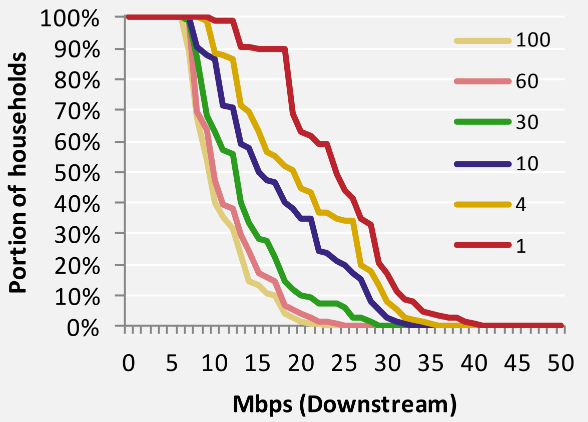

Future scenarios (downstream)

We now turn to the impact of varying key assumptions about the future, first looking at downstream demand. Figure 28 shows the impact of varying various 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.[147]

[147] Note the partial exception for Mobile OS downloads. As expected time to download is reduced to 15 mins, in some cases the required bandwidth *drops*. This counter-intuitive result is because such downloads are a rare event, and as the download speed is increased, in some cases the average downloading time drops to less than 4 minutes per month. Consequently this high bandwidth time drops out of the '4 minutes excluded' picture of demand

Figure 28: 2023 Downstream download and surfing scenarios (4 mins excluded basis)

Assumption	Base	Sensitivity
Mobile OS download time	15 mins	10 mins
Video content DL (as portion of real time)	25%	10%
Web page load time	3 secs	2 secs

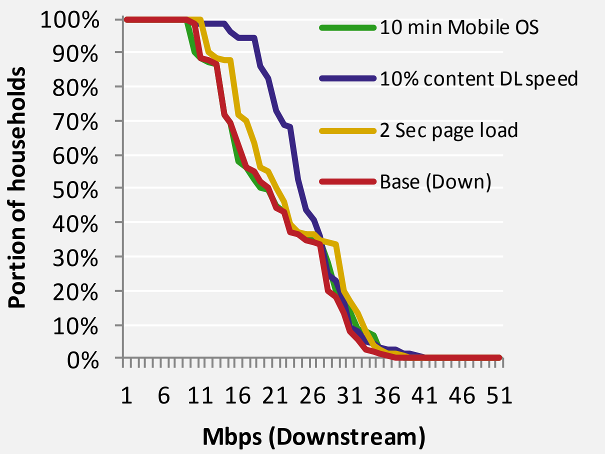

While the various more aggressive assumptions mostly make relatively modest differences to the picture of bandwidth demand, the video content download speed is a partial exception. This change increases the median speed requirement from 19 to 23 Mbps. This is because it forces a significant increase in required bandwidth for content downloading, an activity that is (by 2023) both widespread and frequent.

Figure 29 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 2.5 minutes, then 16% of households require 83 Mbps (and tightening this further would quickly take demand above 100 Mbps).

Figure 29: 2023 Downstream software DL expectation scenarios (4 mins excluded basis)

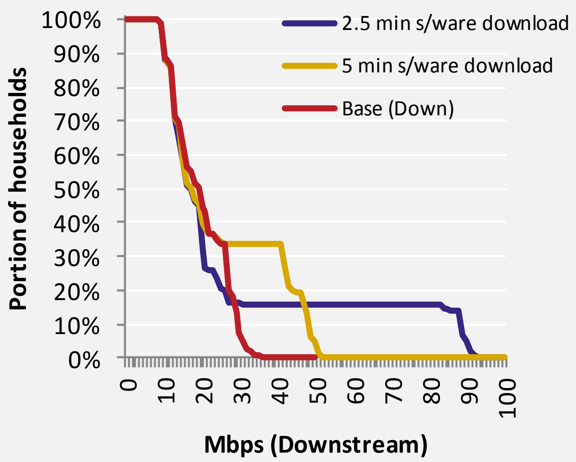

Note however that the downloading of large software files is a comparatively rare event. Consequently, for some households tightening the download time expectation means that downloading is reduced to less than 4 minutes per month, and so is excluded from the demand picture shown here. (Hence the reduction in demand for households in the 16-35% bracket as expectations tighten from 5 to 2.5 minutes).

Figure 30 considers scenarios related to usage overlap and IPTV. 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 moderate. This is because even at four hours, there is already a large amount of overlapping usage.

An increase in 4K TV adoption to 40% also has a moderate impact, at least in the middle range of households – in the base case, the top end households are mostly those that have 4K TV anyway, and thus increased penetration can have little impact there. An increase in starting (2013) bandwidth for 4K TV from 30 to 40 Mbps again has only moderate impact – in part because compression narrows the gap between these two starting assumptions by 2023.

Figure 30: 2023 Downstream IPTV and busy hour scenarios (4 mins excluded basis)

Assumption	Base	Sensitivity
Busy hours (containing 50% of traffic)	4	2
Initial 4K bandwidth (Mbps)	30	40
4K TV penetration (2023)	19%	40%
Annual improvement in video compression	9%	6%

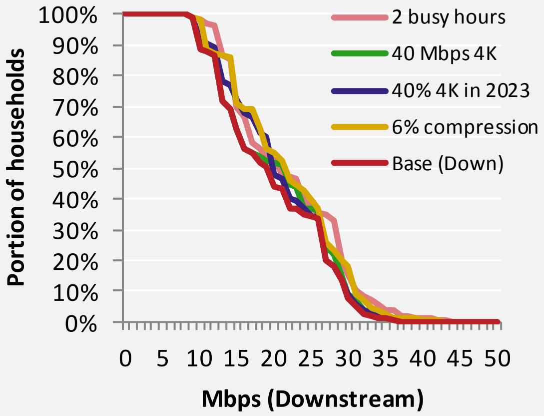

A change to slower improvements in bandwidth compression affects all households, and particularly those at the top end. Video compression drives required bandwidth for IPTV, for content downloads, for video calls and various other applications. Thus this is a key assumption, and particularly in the most demanding households, whose peak demand is in part driven by multiple video streams.

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 31). Note that this scenario is not an 'upper bound' of potential demand – it is possible to imagine users being even more impatient for OS and content downloads, for example.

Figure 31: 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	19%	40%
Busy hours	4	2
Video compr'n CAGR	9%	6%
Initial 4K bandwidth	30 Mbps	40 Mbps

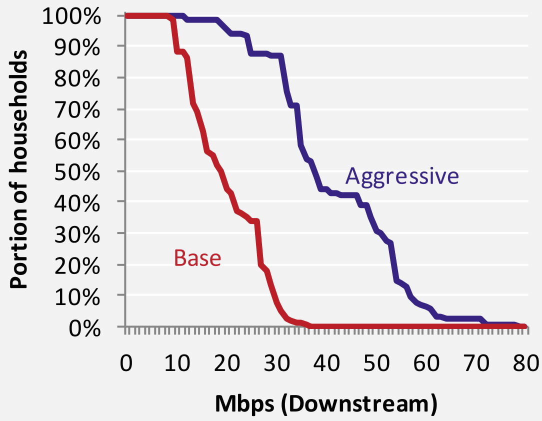

In this scenario median '4 minutes excluded' demand is 38 Mbps, and the top 1% of households have demand of 71 Mbps or more.

Future scenarios (upstream)

We have also considered sensitivities for upstream demand (Figure 32).

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 32: 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]

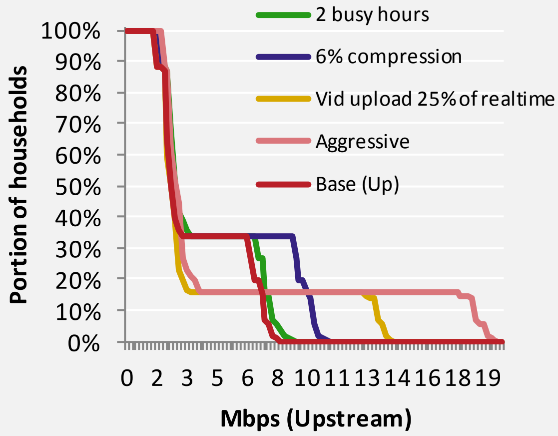

If user expectations of such uploads are that they will take place in 25% of realtime (rather than 50%), this has a dramatic impact. The top 16% of households will require 12.6 Mbps or more. (At these speeds, for the next tier down the time spent uploading drops below the 4 minutes per month threshold, and so 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 the top 16% needing 17.6 Mbps or more.

Comparison to other forecasts

The base case of the model suggests that for all but 1% of households, 35 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 71 Mbps.

These figures are lower than those sometimes informally offered for likely future bandwidth demand. 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.[148] As a consequence, increases in capacity led quickly to increases in demand. To take a simple example, IPTV has been being expected since the days of 28.8 kbps modems.[149] 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.[150]

[148] Certainly for some households this remains so today, as the model suggests
[149] Edmund DeJesus, "How the Internet will Replace Broadcasting", *Byte*, February 1996
[150] FTTH Council Europe, *Creating a brighter future*, 20 February 2013

In allowing for the possibility of a divergence of capacity and demand, we are implicitly rejecting the view that 'build it and they 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 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.[151] (By contrast to the clear benefits of moving from the rudimentary video of a decade ago to today's IPTV).

Equally, while the model anticipates 'conscious usage' of the internet[153] in the home rising to a substantial 4.5 hours per day per individual in 2023 from 2.1 today, this is a slower rate of growth than in the past. Given that people spend on average almost 20 hours per day either out-of-home, sleeping, cooking or doing household chores and admin, 4.5 hours of internet usage represents a significant percentage of the time that could conceivably be spent online. Time spent on entertainment and personal interests totals only 3.82 hours, and this includes linear TV, which is likely to remain significant.

Figure 33: Average time spent by type of activity per individual per day (hours)[152]

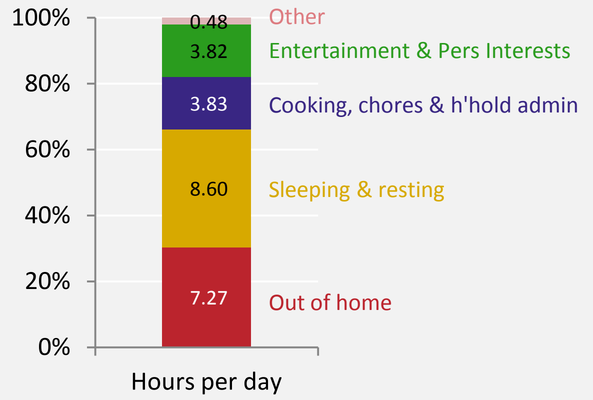

However, there are certain dimensions where we are some way from human limits, in particular impatience. As we have seen, user expectations of speed of content and software downloads is an important sensitivity.[154] Those unwilling to wait more than seconds for very large files to download may demand very high bandwidths.

[151] 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

[152] Office of National Statistics, *Time of Use Survey 2005* [latest available]

[153] Excluding background activities such as file downloads. Note that the 4.5 hours is the sum of individual uses. Since these may be multitasked, the elapsed time will be less

[154] See page 53

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 UK traffic figures are instructive. Between 2012 and 2017, they forecast that consumer fixed internet traffic will grow by 14% annually. However, this growth is almost entirely due to an increase in video, which they see growing at 19% annually from an already high base. They forecast a growth rate of 11% for their 'web and other' category, and see filesharing traffic *falling* by 2% annually. (These are their three key categories of traffic).

Figure 34: Cisco VNI forecast of UK consumer fixed internet traffic[155]

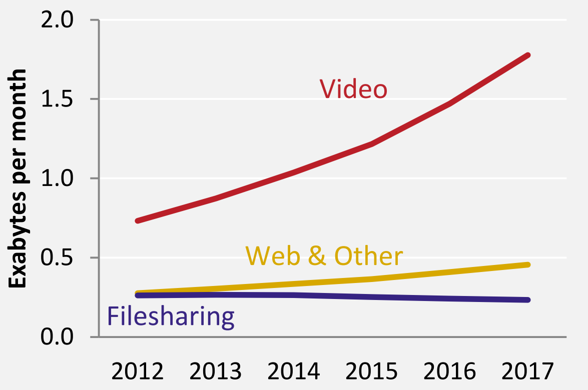

This variation in growth across these categories is particularly relevant in the context of understanding bandwidth. Video is a very significant driver of traffic, but (in common with all streaming applications) is relatively less important in bandwidth terms. This is because its traffic (and traffic 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 IPTV 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.[156] This unsourced view compares to Skype's requirement of 1.5 Mbps for HD video calls.[157] (Even group video with 7 parties – presumably a *very* rare event at home – only requires 8 Mbps

---

[155] Derived from Cisco VNI Forecast widget, set for UK consumer fixed internet traffic [accessed 1 October 2013]. See also Cisco VNI Forecast Highlights for the UK

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

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

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.[158] 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.

[158] 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

We encourage others to develop their own analyses along these lines, to either validate our work or show where it is in error.

Finally, we note the limits of the model's scope – it is a forecast of bandwidth demand. 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 bandwidth service or what products ISPs will seek to offer consumers over time.

8. Assumption summary

			2013	2023	Notes
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%			'Child' is ages 0-15
Household usage variation					
	Share of HH	*Share of Usage*			
Low	40%	20%			
Medium	40%	40%			
High	20%	40%			
Video bit rates (Mbps)					
SD			2	1	
HD			5	2	
4K			30	12	
4K TV penetration			0%	19%	
Video compression improvement (per year)		9%			
Video usage (Hours/adult/day)					
IPTV			0.12	1.09	
YouTube			0.17	0.58	
Content downloads					
Hours of video downloaded per month			3.00	7.77	
DL time as fraction of real time			50%	25%	
Portion of users downloading content			11%	50%	
Mobile OS downloads					
Download time expectation (mins)			30	15	
OS Size (MB)			770	1,997	
Software downloads					
Download time expectation (mins)			60	10	
File Size (GB)			7.0	15.1	
Pre- & post-loading benefit			0%	90%	
Using households		20%			Games consoles are primary driver of s/ware downloads
Downloads per year			0	0	
Video calling					
Usage (mins/adult/month)			54	300	
Bandwidth (Mbps)			0.50	0.19	
Compression improvement		9%			
Video uploads					
Using households		20%			
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)			1.80	2.42	
Average page weight (MB)			1.01	3.47	
Max 'above the fold' page weight (MB)			2.00	2.00	
Time on page (seconds)			53	53	

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 1/8th of a Gigabyte (GB) and equals 1,024 Megabits (Mb)
HD	High Definition, a higher resolution set of video than Standard Definition. HD TV can be transmitted in various formats, the most popular including 720p, 1080p and 1080i

IPTV	Internet protocol television, a method of distribution for professional television / video content 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