



Beyond fast »

How the speed of residential internet access
will develop between now and 2022

Commissioned by:

NLkabel and Cable Europe

Project:

2016.016

Publication number:

2016.016-1617 v4.2

Published:

Utrecht, 10 October 2016

Authors:

Tommy van der Vorst MSc
Reg Brennenraedts MSc MBA
Menno Driesse MSc
Dr. ir. ing. Rudi Bekkers

Management summary

The central question in this study is “How will upload and download speed demand have developed by 2022 in the European market for residential internet access?”. Our research focuses on consumers in the Netherlands and other West European countries with highly developed broadband markets. This study is a follow-up to a study performed in 2014 by Dialogic and Eindhoven University of Technology (TU/e). [4]

Future demand

We project that for the period of 2016-2022, annual growth will be equal to 40.5% for downstream and 44.1% for upstream traffic volume. The growth of existing services alone is responsible for a CAGR of 36.6% (upstream) and 31.6% (downstream) respectively.

The volume growth estimates equal our earlier estimate for the period 2013-2020. We find that our earlier predictions match observed traffic volumes. It appears however that the distribution of downstream and upstream traffic over the course of a day both has become more uniform compared to our measurements in 2014. In 2016, consumers on average take 16.5 hours to upload 80% and 14.5 hours to download 80% of their daily traffic volume, which is longer than observed in 2014. As a result, while volumes grow, the required speeds grow less than expected earlier.

Figure 1 shows that an average subscription will have a sufficient provisioned downstream speed of about 355 Mbit/s in 2020 (compared to 44 Mbit/s in 2016) and an average sufficient provisioned upstream speed of 37 Mbit/s. Table 1 below shows the sufficient speeds for different user groups.




Figure 1. Estimated development of the average sufficient provisioned speed of subscriptions

Table 1. Forecasted average sufficient provisioned speeds (in Mbit/s) for different user groups

		2016	2017	2018	2019	2020	2021	2022
Power users	Up	68	96	135	191	269	379	548
	Down	360	486	656	886	1199	1623	2265
* Note power users:		The estimations for the sufficient provisioned speeds for power users are based on a different method in which traffic for peer-to-peer is modelled to be supply-driven rather than demand-driven. This means that the power users will always maximally utilize the provisioned bandwidth.						
Innovators	Up	9	12	17	25	35	49	71
	Down	111	149	201	271	366	495	693
Mainstream users	Up	1	1	2	2	3	4	6
	Down	17	22	31	42	57	77	109
Laggards	Up	0	0	0	1	1	1	2
	Down	2	3	4	5	7	9	17
All users	Up	4	6	8	12	17	25	37
	Down	44	62	87	122	172	242	355

Traffic volume growth is primarily driven by online video and music services (either streaming or peer-to-peer). Figure 2 below shows the estimated demand for downstream traffic volume broken down by service category.




Figure 2. Forecasted average daily downstream traffic volume per residential subscription

Table of contents

1	Introduction	7
1.1	Research questions	7
1.2	Research methodology	7
1.3	About the researchers	10
1.4	Reading guide	10
2	Current demand	11
2.1	Total traffic demand	11
2.2	Demand by service	12
2.3	Demand for speed	14
2.4	Differences between user groups	16
3	Demand growth	19
3.1	Aggregate traffic volume growth	19
3.2	Growth by service	21
3.3	Future revolutionary services	24
3.4	Differences between user groups	25
4	Future demand	27
4.1	Aggregate demand	27
4.2	Demand by service	27
4.3	Demand for speed	29
4.4	Differences between user groups	30
	Appendix A. Methodology	35
	Overview	35
	Modelling aggregate demand	36
	Modelling demand by service	36
	Modelling differences between user groups	37
	Estimating speed demand for a given level of traffic volume	39
	Modelling demand growth	39
	Limitations	41
	References	43

1 Introduction

Ever since the introduction of the internet, the demand for speed on residential internet access connections has continued to grow. Even though home access connections provide generous amounts of bandwidth nowadays, records are still being broken. The end of growth in bandwidth demand seems to be nowhere in sight.

In 2014, Dialogic and Eindhoven University of Technology (TU/e) performed a study on the development of demand for residential (consumer) internet speeds in the period between 2014 and 2020. [4] The report now in front of you provides an updated projection for the period between 2016 and 2022. The update was compiled by using the methodology employed in 2014, but with new data obtained from operators and from literature. Whenever appropriate, we compare the findings and projections from 2014 with our updated conclusions.

1.1 Research questions

The central question in this study is “*How will upload and download speed demand have developed by 2022 in the European market for residential internet access?*” In order to provide the answer, we first need to address the following set of sub-questions:

1. To what extent do currently available applications contribute to internet traffic?¹
2. To what extent has the need for traffic of currently available applications changed in recent years?
3. To what extent will consumers use other applications with a high demand for internet traffic by 2022?
4. Which upstream and downstream subscription speeds will be sufficient for future demand?

Our research focuses on consumers in the Netherlands and other West European countries with highly developed broadband markets.

It is essential to decide on a time horizon in order to formulate conclusions relevant for defining policy as well as strategy. As the broadband market is highly dynamic, we have chosen a time horizon of seven years (up to and including 2022). Policy makers generally have a slightly longer time horizon with respect to digitalisation, generally up to 2025. We feel that due to the uncertainty with respect to online service development and innovation we cannot make estimates further than seven years into the future. This is substantiated by our analysis of the sensitivity of our model.

1.2 Research methodology

The research methodology employed to answer these questions is described in more detail in Appendix A. The methodology is equal to the methodology used in our 2014 report, except for two aspects. The first is that in this updated research, we use new and updated data sets, some of which were obtained from other sources than in our 2014 report. Specifically, we

¹ With respect to both upstream and downstream traffic, we distinguish between traffic *volume* and (peak) *speeds*. Throughout this document we use upstream or downstream *traffic* to indicate volume.

have now primarily relied on data obtained from two rather than three ISPs (labelled 'ISP A' and 'ISP B' in this report). Second, we did not perform an interview round with experts in this edition.

1.2.1 Traffic and speed

In this report, we frequently refer to the terms 'traffic' and 'speed'. Using a service leads to a certain amount of data that needs to be transferred between a household and the service provider over the access connection. The data has a certain size and needs to be transferred in a certain time period. With *traffic*, we refer to the total amount of data to be transferred in a given time period (usually measured as megabytes per day or month).

The *speed* of a connection refers to the amount of traffic it can transfer in a given time period (usually measured in megabits per second). The speed of a subscription as advertised by ISPs is referred to as the *provisioned speed* and refers to the maximum speed attainable over the connection. The term *bandwidth* is similar and usually indicates the net availability of speed at a particular moment. In this report, we consistently take overhead into account and exclusively present results in terms of *provisioned speed*, in order to improve comparability of our results with figures quoted by the ISPs.

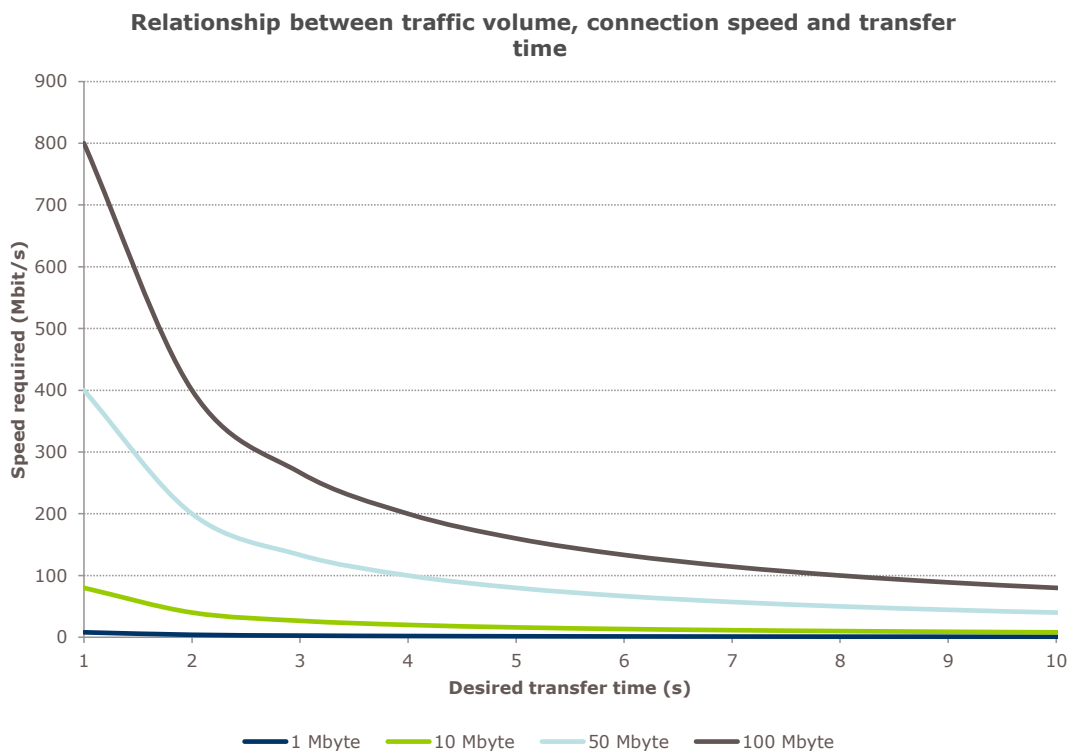


Figure 3. The relationship between traffic volume, connection speed and desired transfer time

Because traffic is unevenly distributed over time (e.g. most users will not use their internet connection as much at night as they do during the day), the peak speed required by a household will be (many times) higher than the *average* amount of traffic per second. To illustrate: while a connection with a downstream speed of 0.1 Mbit/s is sufficient to download a 1 GB movie in a day, the connection does not meet user *demand* if they want to watch the movie right away (which requires the 1 GB movie to be transferred in much less time, e.g. half an hour). Figure 3 illustrates the relationship between speed and traffic by showing the speeds required to transfer various amounts of data in various time periods.

In this report, speed is always expressed in megabits per second (Mbit/s) whereas traffic is expressed in megabytes per day or month (Mbyte/day, Mbyte/month).

1.2.2 Modelling approach

The speed demand estimation model developed in this study is calibrated using actual measurements and estimates of the traffic volume generated by different service categories and by different user groups (see diagram in Figure 4).

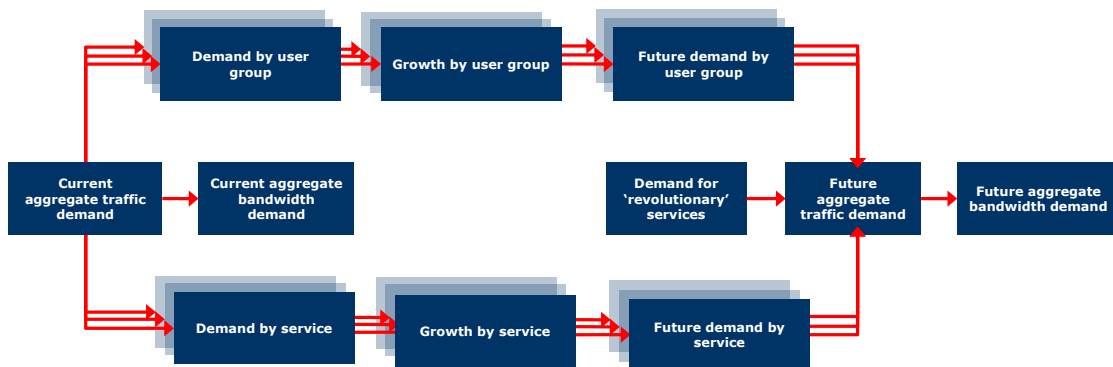


Figure 4. Schematic overview of the demand model

In order to estimate the total traffic volume of residential subscribers, we started with the current aggregate traffic measurements obtained from service providers. To estimate future demand, we broke down these statistics (both historic and current) into two categories: demand by user group and demand by service. Subsequently, we applied growth factors (as a result of adoption as well as intensity growth) to each service and user group. Adding up the individual services together then leads to an aggregate of the total traffic required.

Because we were not only interested in traffic but also speed demand, the final step was to estimate speed demand based on current traffic demand. This was done for both current and future demand.

A more detailed description of the modelling approach can be found in Appendix A.

1.2.3 Scope of the model

For the purposes of this study, we made several key choices regarding the model's scope and boundaries. The aim is to estimate speed demand over a time horizon of seven years. The unit of analysis is household connections. Although we do not model individual households, we aim to distinguish different groups as well as indicate the distribution of demand over households. Note that we thus only model the use of business applications for connections with a consumer subscription; small offices or home offices (SOHO) using a non-consumer connection are excluded.

Geographically speaking, we are interested in those West-European countries where residential broadband connections are commonplace. We therefore assume that the *exogenic* growth would be negligible, meaning that we do not expect a significant impact on the demand resulting from first-time internet users.

We model demand for speed rather than availability of certain internet speeds, i.e. traffic consumption is driven by demand rather than supply. We assume that the available bandwidth is secondary to the demand: this means that operators will use the expected demand as a guideline for dimensioning their network.

A more detailed description of the way the scope is implemented in our model can be found in Appendix A.

Traffic and speed demand for specific purposes / sectors

This study is concerned with *generic* demand for broadband at home. In specific sectors and for specific purposes, broadband demand may differ significantly from this generic demand. The results from this study can however be used in these specific cases as well as a basis for modelling broadband demand.

Consider, for instance, the educational sector. Following a similar approach as in the presented results, Dialogic provided (among others) the expertise behind the '*Handreiking externe connectiviteit*' of Kennisnet [8]. This guide supports elementary and secondary schools in their selection of (internet)connectivity. It echoes our expected growth rate in required provisioned speeds, namely a doubling within three years. At schools, generic demand is found for instance in provisioning public Wi-Fi to students. In addition, schools have extra requirements regarding the reliability of a link (e.g. for facilitating digital exams).

1.3 About the researchers

This study was conducted by Dialogic *innovatie & interactie* and the Eindhoven University of Technology (TU/e). Dialogic is a research consultancy based in the Netherlands, focused on innovation and specialised in telecommunications. In the past fifteen years Dialogic has conducted studies for many clients in the public and private domain, in both the Netherlands and internationally, among which are ACM (the Dutch national regulatory authority for telecommunications), the European Commission, ITU, Dutch Ministry of Economic Affairs, most Dutch provinces.

Eindhoven University of Technology (TU/e) is a large research and educational institute based in the south of the Netherlands. Dialogic maintains close connections with the worldwide academic community, and frequently collaborates with Eindhoven University of Technology on various telecommunications-related projects.

1.4 Reading guide

This report largely follows the same structure as our 2014 report, except for the description of the methodology, which has been moved to a separate annex. In chapter 2, we describe the current demand for internet speed. In chapter 3, we describe the developments that will lead to demand growth over the period until 2022. In chapter 4, we present the projected demand based on the findings from the two preceding chapters.

2 Current demand

We now turn to modelling the current demand for speed. Starting with the aggregate demand (which concerns all internet users), we provide further details by distinguishing various service categories. We then make the step from traffic volume demand to speed demand by looking at the time during which traffic is generated. This enables us to finally estimate the current speed demand for the various user groups.

2.1 Total traffic demand

In our 2014 report, we predicted that an average internet user would demand 41,933 Mbyte per month of traffic (upstream and downstream combined). We projected this demand to grow to 59,165 Mbyte/month in 2015, and to 83,586 in 2016. Measurements on the networks of ISP A² and ISP B³ allow us to verify these projections. Figure 5 shows the predicted and measured traffic average monthly traffic volume. As the figure shows, our estimates quite closely match the traffic demand of the average user on the networks of both ISP A and ISP B. Note that as our model provides estimates at the level of a year, it does not account for variations between quarters *within* a year (hence, the measurements 'swing around' our predictions in Figure 5).

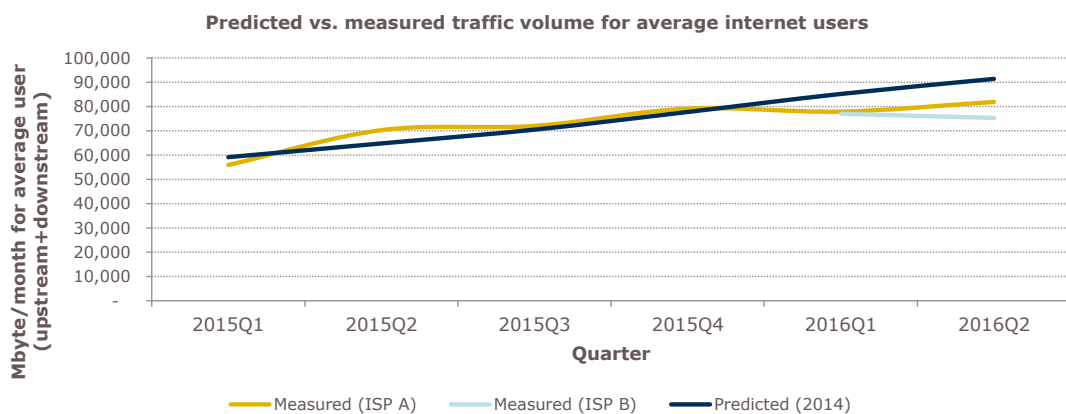


Figure 5 Predicted [4] and measured average monthly traffic volume

Other sources further strengthen our conclusion that the model provides accurate results. Table 2 shows additional data points, obtained from various other sources.⁴

² ISP A provided us with aggregate average traffic volumes distinguished by subscription category. We translated these figures to an average using the distribution over user groups shown in Table 4. The traffic is measured at the CPE level and does not include IPTV traffic.

³ On the network of ISP B, traffic measurements were taken from a core router that carries traffic for consumer and SME subscriptions over both FttH as well as cable. The speed limits associated with subscriptions on each infrastructure are equal. The aggregate traffic includes IPTV traffic, but as this traffic is multicast, it is not multiplied for each subscriber and hence does not amount to a significant share of traffic.

⁴ Unfortunately, most of the sources mentioned in our 2014 have stopped reporting totals and per-household averages, or have merged upstream and downstream traffic figures. For comparability we chose to display total traffic volume per month in Table 2.

Table 2. Average total monthly volume of traffic per household, according to various sources. The figures in grey were also mentioned in our 2014 report.

Source	Period	Location	Total traffic volume (Mbyte/month)
Sandvine [14]	2013 H1	Western Europe	13,400
Sandvine [14]	2013 H2	Western Europe	17,400
ISP A [4]	2013-11	The Netherlands	7,466
ISP B [4]	2013-11	The Netherlands	2,655
Cisco VNI [3]	2014	Western Europe	38,800
ISP A	2015-Q1	The Netherlands	55,955
ISP A	2016 Q1	The Netherlands	77,924
ISP B	2016-Q1	The Netherlands	76,986
ISP B	2016-Q2	The Netherlands	75,315
Cisco VNI [3]	2015 (predicted)	Western Europe	46,661
Cisco VNI [3]	2016 (predicted)	Western Europe	56,115
Cisco VNI [3]	2017 (predicted)	Western Europe	67,484
Cisco VNI [3]	2018 (predicted)	Western Europe	81,157

Note that Cisco VNI predicts a rather conservative CAGR (compound annual growth rate) of 20% for Western Europe. This is especially interesting given the fact that the approach taken by Cisco VNI in coming to these estimates is highly similar to our approach. We suspect that the difference in outcome is therefore the result of sampling bias. The Cisco numbers represent the Western European region, which includes countries with a suboptimal infrastructure. Our measurements and modelling primarily concern the Dutch market, which is known for its high penetration of high speed and high quality broadband infrastructure and ICT services.

2.2 Demand by service

Allocating the demand to the various internet-based services plays a central role in our prediction model. In our 2014 study we distinguished 14 services, ranging from consultative web browsing and e-mail to online gaming and peer-to-peer file sharing. By combining top-down and bottom-up figures as well as approximation we came to a distribution of the demand per service. Subsequently, we were able to forecast future demand of the individual services. For this update of the study, we recalibrate our earlier predictions using recent observations. From this exercise we can determine whether we need to adjust our modelled distribution for the new time frame (now-2022).

The Global Internet Phenomena reports by Sandvine present peak period traffic composition on the top 5 services each year, which provides much of the required insight. Figure 6 and Figure 7 show the results of their bi-annual reports since 2012. Unfortunately, Sandvine no longer presents total traffic volume in their reports, so we are only able to discuss the trends in the ratios.

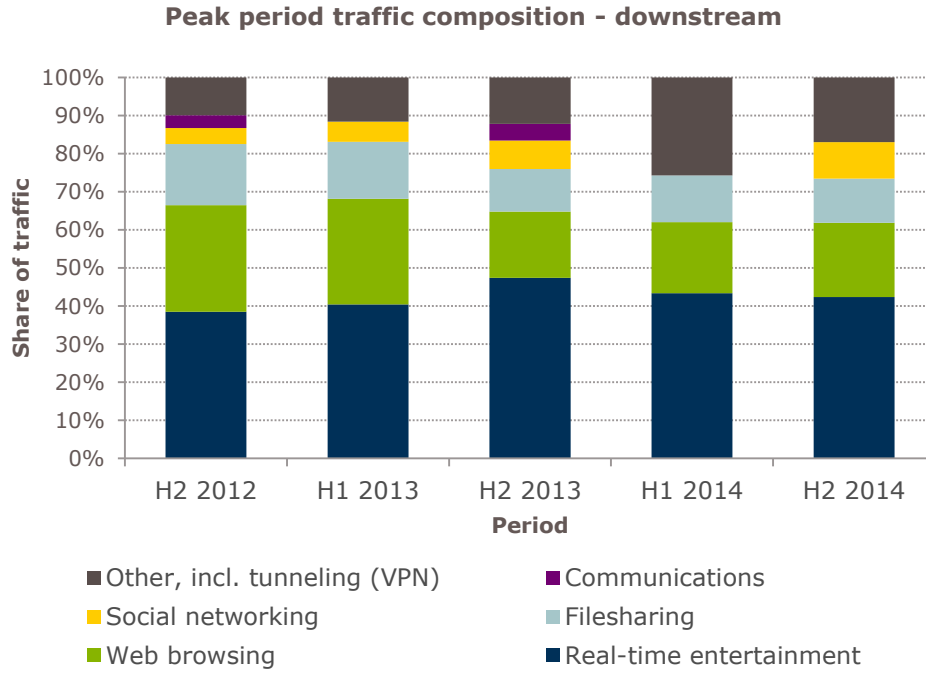


Figure 6. Sandvine measurements of the distribution of downstream traffic volume by type [14]⁵

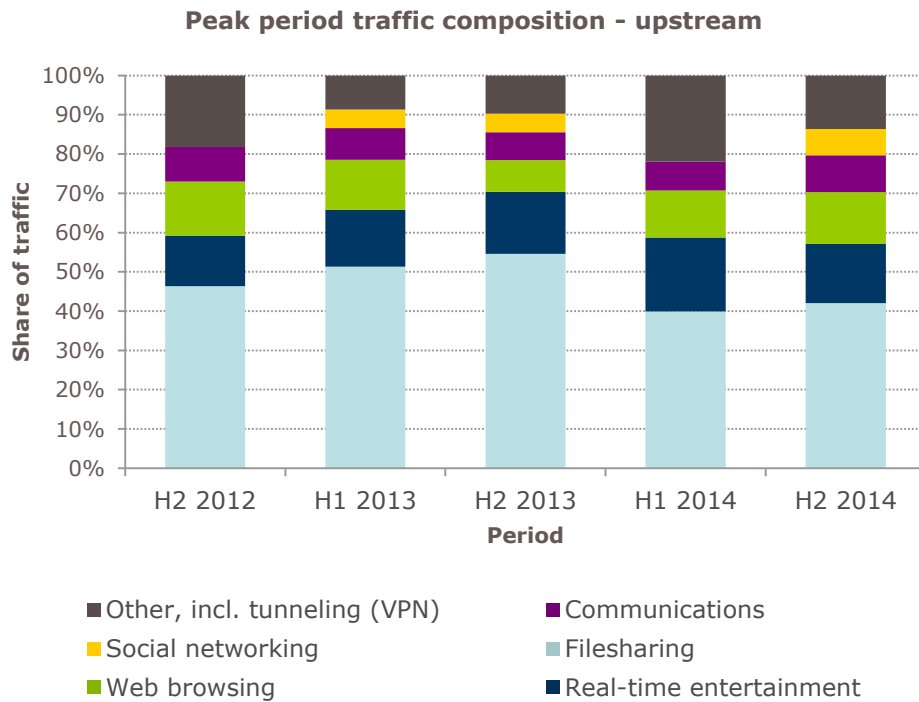


Figure 7. Sandvine measurements of the distribution of upstream traffic volume by type [14]

⁵ We added *marketplaces* (eBay, Amazon, etc.) and *tunnelling* (VPN) to *other* (all services outside of top 5) for the periods where they were presented separately.

Over time, there are some clear variations of the relative share of each service. The total downstream share of online video and web browsing is however fairly stable over the entire presented period. Since they together represent over half of the total downstream volume, we conclude there is limited need to make major adjustments in this section.

For upstream however, we observe a strong decrease in the share of peer-to-peer file sharing in upstream traffic volume. The absence of the total volume figures in the latest report limits us to further examine this trend. The broader uptake of music and video streaming service could explain a decrease in the demand for P2P-services, since it (to some extent) substitutes piracy related torrent/P2P traffic. Additionally, the growth of traffic makes the (non-elastic) group of peer-to-peer users become relatively smaller.

Foregoing analysis enabled us to recalibrate our 2016 volume figures. The new data led to minor adjustments in the upstream volume distributions, but only to a limited extent. Table 3 shows the results. The output of our 2014 model proved to be highly accurate when compared to actual measurements on the network of ISP A.

Table 3. Services and the average amount of traffic generated per subscription

Service	Average daily volume (Mbyte) 2013		Average daily volume (Mbyte) 2016	
	Up	Down	Up	Down
Consultative web browsing	10.4	99.4	22.9	219.0
E-mail	0.3	5.0	0.4	6.7
Social media / Web 2.0	3.9	18.9	10.4	51.2
Remote backup	25.0	0.0	78.1	-
Conversational applications	6.6	6.6	10.5	10.5
Remote workplace	17.4	174.5	21.6	215.5
File downloads	0.0	10.4	-	14.9
Online video and music	83.5	283.6	227.5	791.9
Online gaming	0.0	0.1	0.0	0.2
Personal cloud storage	1.2	3.0	2.5	5.9
<i>Other services</i>	16.5	82.0	31.4	156.4
<i>Overhead</i>	76.6	50.7	167.5	73.6
Net total	241.5	734.2	573	1,546
Net total (monthly)	7,365.8	22,393.2	17,473	47,150

2.3 Demand for speed

In our 2014 report, we analysed in depth the relationship between traffic volume and the connection speed required by consumers. [4] Due to the fact that consumers do not use their internet connection equally over the course of a day, the required connection speed is not the same over the course of a day, but rather exhibits peaks. From data analysed in 2014, we concluded that downstream traffic increases gradually over the course of a day and peaks at around 7:30 pm. A similar pattern was observed for upstream traffic, although the increase is much less (that is, the upstream traffic volume stays more constant over the course of a day). The traffic distribution is the determining factor to answer the question what the minimum acceptable speed is for a connection, given a certain level of traffic demand.

Using data from ISP A and ISP B, we were able to replicate these findings for 2016. Figure 8 shows the cumulative distribution of traffic over the course of a day, based on measurements on the network of ISP B in 2013 and 2016.⁶

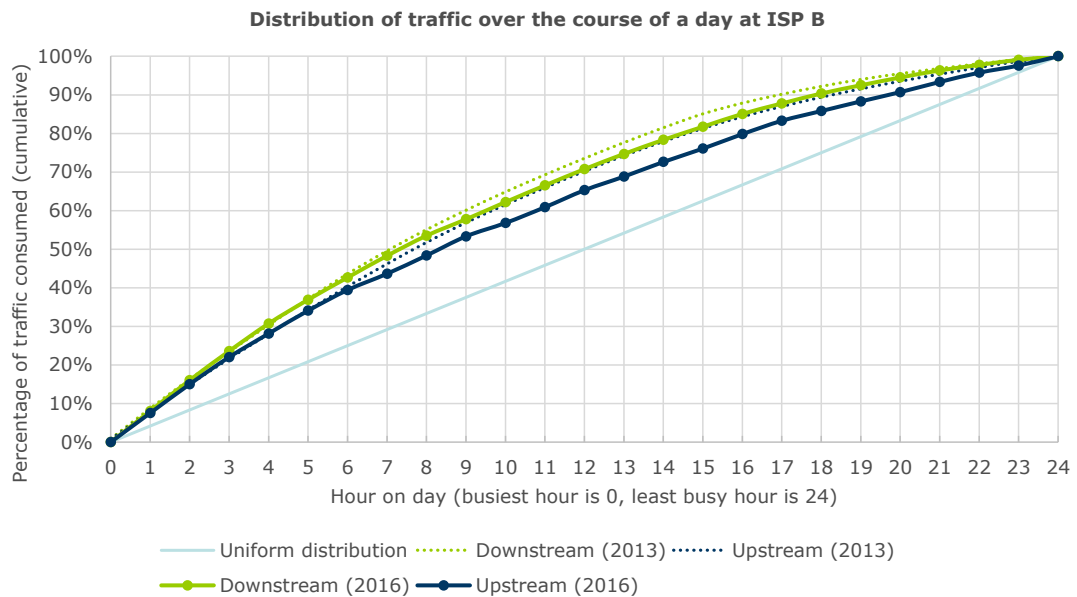


Figure 8. The cumulative proportion of total traffic volume transferred in a day by minutes (measured on the ISP B network in May 2016, as well as the results obtained in 2013 [4])

The distribution of downstream and upstream traffic over the course of a day both have become more uniform compared to our measurements in 2014. [4] In other words, the connection speed required over the course of a day more closely follows the speed that would be required when the consumption of traffic is equally distributed over the day (depicted with the dotted line in Figure 8)

The difference in concentration of upload versus download traffic has however increased. From Figure 8, one can see that in 2016, consumers on average take 16.5 hours to upload 80% of their daily upstream traffic, whereas they take 14.5 hours to download 80% of their daily downstream traffic. In 2016 it takes almost two hours longer to consume 80% of upstream compared to downstream traffic volume, compared to only one hour in 2014.

The result that downstream traffic is less evenly distributed over time than upstream traffic leads to a difference in the minimum connection speed required between upstream and downstream. Given the same daily volume, less bandwidth would be needed upstream than downstream.

⁶ We have also calculated the distribution of traffic based on measurements from ISP A and found that it follows the pattern shown in Figure 8.

Why has traffic become more evenly distributed?

We did not obtain access to data detailed enough to investigate this question at length. We suspect however that the cause for the decreased urgency of traffic is at the application level. The first hypothesis is that existing applications (and their usage) have changed in such a way that traffic has become less urgent. The level of urgency of an application's traffic is for instance impacted by decisions to increase the video buffer size or by implementing background downloading or preloading of content. We have not found specific applications that have implemented such a change and that could have caused a substantial change in aggregate traffic urgency.

The second hypothesis is that new applications with a lower (average) level of traffic urgency have increased in adoption relative to other applications that have more urgent traffic. For instance, streaming video services may have partially substituted traffic of peer-to-peer file sharing. While in file sharing complete movies or even seasons of television shows are downloaded once and watched later, streaming video platforms only download while watching. Therefore, streaming video traffic is much less 'urgent' than is file download traffic.

The difference between upstream and download traffic urgency has also increased. In our 2014 report, we described several explanations for this phenomenon. [4] It is likely that since 2014, these have caused the urgency of downstream versus upstream traffic to drift apart even further.

2.4 Differences between user groups

Following our 2014 report, we assigned traffic shares to different user based on measurements of traffic distribution. Table 4 gives an overview of the groups, their share in the upstream and downstream traffic generated and how the respective shares compare to the total average.

Table 4. User groups, the share of traffic they each generate and the parameters chosen for their adoption curves

	Lag-gards	Main	Innova-tors	Power users
Percentage of subscriptions	20%	60%	18%	2%
Percentage of upstream traffic	1%	10%	44%	45%
Percentage of downstream traffic	1%	29%	52%	18%
Upstream traffic compared to average user	0.1x	0.2x	2.4x	22.5x
Downstream traffic compared to average user	0.1x	0.5x	2.9x	9.0x

Given the groups as shown in Table 4, it is possible to estimate the demand for services as well as total traffic demand in absolute terms for each user group; the results are shown in Table 5.

Table 5. Estimated current demand for traffic volume by different user groups

	Up Mbyte/day	Down Mbyte/day		Up Mbyte/day	Down Mbyte/day
Power users	11,690	14,918	Innovators	1,503	4,788
Consultative web browsing	513	1,960	Consultative web browsing	56	629
Conversational applications	198	79	Conversational applications	22	25
E-mail	10	60	E-mail	1	19
File downloads	-	124	File downloads	-	40
Online gaming	1	2	Online gaming	0	1
Online video and music	4,788	5,754	Online video and music	520	1,847
Other services	750	1,459	Other services	84	476
Personal cloud storage	33	31	Personal cloud storage	0	0
Remote backup	1,381	-	Remote backup	150	-
Remote workplace	453	1,813	Remote workplace	49	582
Social media / Web 2.0	213	418	Social media / Web 2.0	23	134
Overhead	1,587	585	Overhead	421	188
Future revolutionary services	1,765	2,633	Future revolutionary services	177	846
Mainstream users	124	1,507	Laggards	34	75
Consultative web browsing	4	105	Consultative web browsing	1	11
Conversational applications	1	4	Conversational applications	0	0
E-mail	0	3	E-mail	0	0
File downloads	-	7	File downloads	-	0
Online gaming	0	0	Online gaming	0	0
Online video and music	34	282	Online video and music	0	1
Other services	-	43	Other services	19	42
Personal cloud storage	0	0	Personal cloud storage	0	0
Remote backup	9	-	Remote backup	0	-
Remote workplace	3	97	Remote workplace	1	9
Social media / Web 2.0	2	22	Social media / Web 2.0	0	1
Overhead	59	28	Overhead	8	3
Future revolutionary services	11	915	Future revolutionary services	5	7

2.4.1 Provisioned speeds

ISPs provision their connections in such a way that the consumer does not experience limitations during usage. Generally, links are provisioned with much higher peak speeds than would be strictly required following the measured urgency of traffic discussed above. The reason is that while traffic may be evenly spread over the course of hours, traffic generated by many services still exhibits peaks at the sub-second level. While speeds estimated based on the urgency measurements may be 'good enough', user experience is significantly improved when links are provisioned with higher peak speeds.

Table 6. Information on subscriptions from the ISPs whose traffic measurements were used

Subscription	Number of subscriptions	Advertised speed (Mbit/s)	
		Up	Down
ISP A			
Medium	45%-50%	4	40
High	40%-45%	15	150
Extreme	5%-10%	30	300
ISP B			
Basic	8%	5	5
Low	2%	10	10
Medium 1	6%	35	35
Medium 2	59%	40	40
Medium 3	1%	75	75
High 1	20%	100	100
High 2	3%	200	200
Extreme	1%	600	600

By comparing the provisioned speeds to the theoretically required speeds, we have calculated the minimum speed the ISPs would have to provision in order to meet the traffic volume demands discussed earlier. In 2014, we estimated that for an average user, the minimum required provisioned speeds were 21.4 Mbit/s (downstream) and 2.2 Mbit/s (upstream). For 2016, we estimate minimum required speeds of 43.8 Mbit/s (downstream) and 4.0 Mbit/s (upstream).

Table 7 Minimum speeds to provision per household to satisfy 2016 demand for traffic volume

	Up Mbit/s	Down Mbit/s
Power users	68.4	360.3
Innovators	8.7	110.8
Mainstream users	0.8	16.5
Laggards	0.2	2.0
All users	4.0	43.8

Table 7 shows the minimum speeds that need to be provisioned to users in the different groups in order to satisfy the current traffic volume demand for these users. Comparing the table to Table 6 on currently available subscriptions, it appears that the ISPs currently offer subscriptions that map well to the user groups distinguished in our study, except perhaps for the 'power users' who currently cannot obtain the desired upstream capacity from ISP A.

3 Demand growth

In this chapter, we discuss the drivers for demand growth. We start out with an analysis of the development of traffic growth from a high level of aggregation. We then descend to the level of individual households, and analyse the influence of service adoption, intensity growth and new service introduction on their traffic demand.

3.1 Aggregate traffic volume growth

The average total traffic volume (upstream and downstream combined) in 2015 shows a week-on-week growth rate of 0.7% on average, equalling to a year-on-year growth rate of 42.2%. This is shown in Figure 9, where the weekly average traffic volumes as measured on the network of ISP A are expressed as Mbyte/month per subscriber, and compared with a reference growth level of 42%. In our previous report, we predicted a compound aggregate yearly growth rate of 40.5% for downstream and 44.1% for upstream traffic. Note that modelling growth at this level is sensitive for variation between weeks. We can however safely conclude that the growth levels predicted in our 2014 report are accurate as far as the aggregates are concerned.

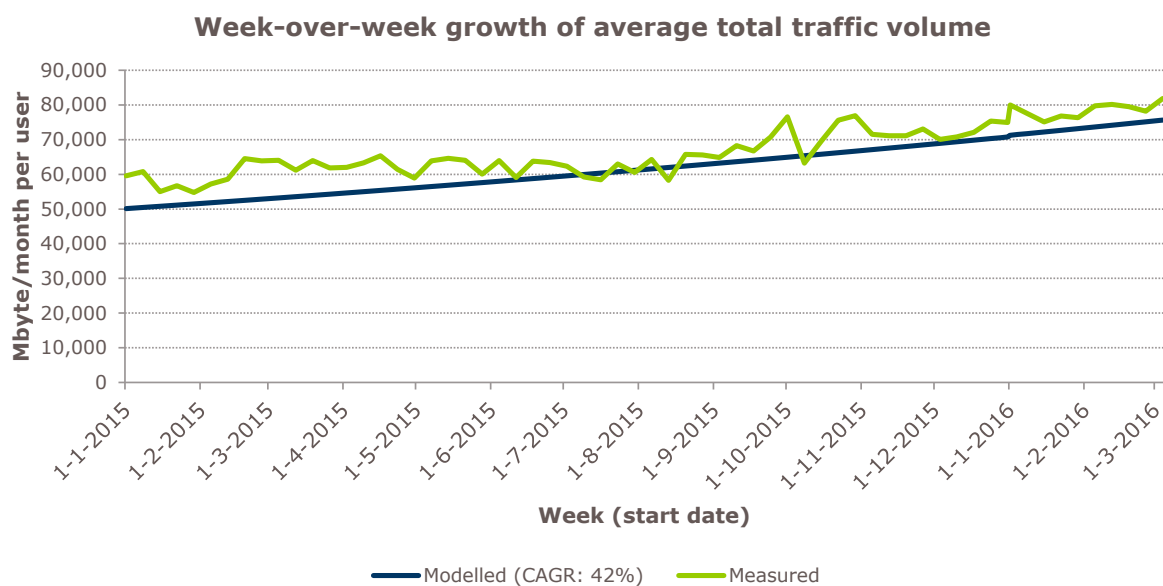


Figure 9 Week-on-week growth of traffic on the network of ISP A, exhibiting a year-over-year growth rate of 42%

In conclusion we project that for the period of 2016-2022, annual growth will be equal to 40.5% for downstream and 44.1% for upstream traffic volume. This equals our earlier projection [4] for the period 2013-2020. This projection is justified by evidence at the ISP level, where measured traffic levels closely follow our earlier projections, as well as the aggregate level, where exponential growth is sustained as well.

Is internet traffic volume still growing exponentially?

Evidence for the hypothesis that growth of internet traffic volume is still following an exponential pattern can also be found in statistics at an even higher level of aggregation. Traffic levels at the Amsterdam Internet Exchange (AMS-IX) provide insight in the total demand for internet traffic at the aggregate level. Figure 10 shows the development of volume of the traffic exchanged over AMS-IX.

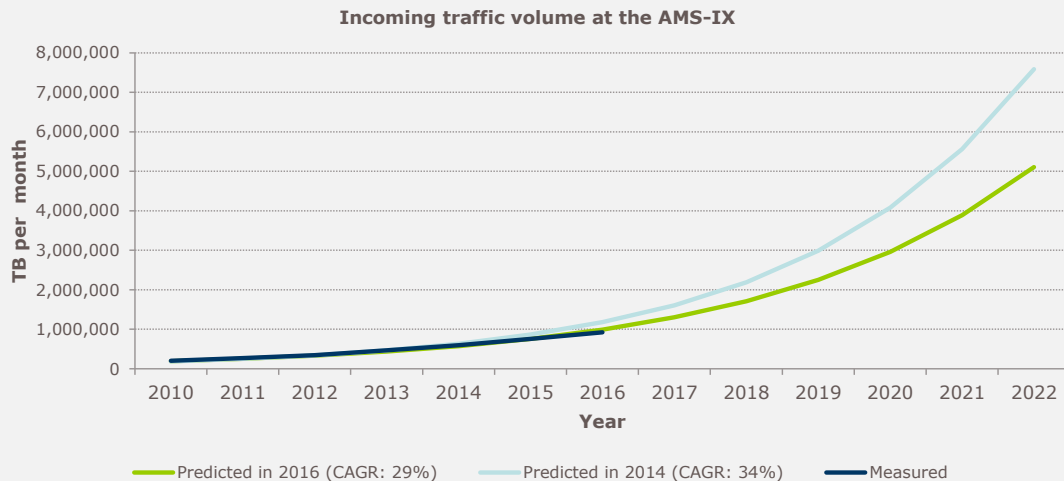


Figure 10 Measured and predicted growth of the average monthly incoming⁷ traffic volume at the AMS-IX [1]

Over the period 2004-2014, a compound annual growth rate of 34% can be observed. Using data over the period 2005-2016, the growth rate observed is 29%. For 2022, we predict that the aggregate traffic flowing through AMS-IX will be 5x that of the volume flowing today, or between 4 and 6 exabyte (million terabyte) per month, on average.

Note that the growth figures at the aggregate level are not comparable to growth figures at the consumer level (which we estimated at 40% for downstream over the period 2014-2020 [4]). A significant portion of traffic from and to consumers does not flow over the AMS-IX. In addition, traffic over the AMS-IX also includes peering traffic between networks that are not necessarily 'eyeball' consumer networks.

While the growth of aggregate traffic appears to have slowed down slightly over the past few years, this also does not necessarily imply that traffic volume demand growth has also slowed down from the consumer point of view. Rather, it may indicate that service providers choose other ways of exchanging traffic with content providers (e.g. using content delivery networks that have points of presence *inside* the ISP network, closer to the edges) or be the result of other technical changes leading to less traffic⁸.

Will internet traffic volume grow exponentially forever?

There are certain physical limits to the speed with which information can be processed and the density at which it can be stored. Therefore, there is a certain absolute upper bound on the maximum volume for internet traffic. There are also reasons to believe that internet traffic volume will reach a plateau before that upper bound is reached. For instance, there may simply be an upper bound to the amount of information a human can process (which is also an argument that is commonly mentioned when discussing resolutions of television screens or digital cameras – at some point, humans simply cannot see improvements anymore).

What can also be observed in such discussions is that often, at a certain point, other characteristics become more prominent differentiators. Digital camera manufacturers today advertise the quality of their lenses or image processing algorithms, or tout additional features, rather than the number of megapixels on the sensor. Subsequently the growth in the number of megapixels seems to have stalled. It is possible that at a certain point, the current focus on speed for internet connections may also shift towards attributes such as latency, jitter, reliability and customer service. An example of the latter is Dutch ISP XS4ALL, which is a premium brand offered by KPN and advertised as having a very good customer service desk (KPN also offers internet service through their Telfort brand, which is advertised as being very low cost).

The focus may also shift towards latency and jitter because of increasing importance of certain applications that require low latency and jitter. Currently, there are high expectations for virtual reality. In order to deliver a high-quality virtual reality experience, low latency is essential to (literally) prevent headaches. Another application that expects low latency and jitter are self-driving cars.

3.2 Growth by service

3.2.1 Service adoption growth

In our 2014 report, we modelled adoption curves for different categories of services. In 2014, many of the service categories were already at or close to 100% adoption. The service categories that we expected to experience the highest growth in adoption were remote backup, online video, and personal cloud storage. Compared with the situation in 2014, we do not see significant reasons to alter the estimate for these, or any of the other service categories. The reason for this is twofold. First of all, we see no support for the hypothesis that service adoption has accelerated or slowed down in the relatively short period of two years since our last research. In addition, none of the service categories has experienced the introduction of a new, disruptive and traffic-consuming service that would skew the results. The fact that the growth modelled in our earlier report quite closely matches the traffic volumes observed is additional evidence to this conclusion.

Table 8 gives an overview of the modelled adoption for each service category.

⁷ As an Internet Exchange primarily exchanges traffic between networks and does not consume or produce traffic, the total volume of incoming traffic is in theory equal to the total volume of outgoing traffic. In practice, the totals differ slightly due to technical reasons (packet loss, disruptions, et cetera).

⁸ For example, more efficient use of multicast and broadcast protocols would lead to lower growth of traffic volume at the aggregate level, but equal or higher growth at the edges.

Table 8. Age and modelled adoption of the service categories distinguished in our analysis

Service	First consumer implementation	Year introduced	Current age (years)	Modeled adoption
Consultative web browsing	WWW	1989	27	100%
E-mail	Sendmail	1983	33	100%
Social media / Web 2.0	Geocities	1994	22	99%
Remote backup	Carbonite	2004	12	43%
Conversational applications	NetMeeting	1995	21	98%
Online video	YouTube	2005	11	32%
Remote workplace	Citrix	1995	21	98%
Online music	MP3.com	1997	19	96%
File downloads	WWW	1989	27	100%
Online gaming	Ultima Online	1997	19	96%
Peer-to-peer file sharing	Napster	1999	17	
Personal cloud storage	Dropbox	2007	9	15%

What about peer-to-peer file sharing?

In our 2014 report, we distinguished peer-to-peer file sharing, online video and online music as three separate service categories. In this edition, we have only included a single category for online video and music, which is calculated as the sum of the former three separate categories.

Adoption of peer to peer services is primarily influenced by the risk posed by viruses or spyware and, to a lesser extent, the penalties for downloading or sharing illegal content. In addition, the introduction of legal alternatives (and their adoption) has a negative effect on the adoption of peer-to-peer services. Indeed, if legal alternatives become successful, the adoption of peer-to-peer may decrease before it has reached its maximum adoption potential. Whether a 'churn' from illegal peer-to-peer towards legal distribution is occurring or will occur in the future is still a topic of debate.⁹ Figure 6 does seem to suggest such a churn is in progress. Additionally we see renewed interest from the content industry to discourage peer-to-peer distribution of content.

The primary reason for merging the categories is that based on the data available to us, we cannot conclusively estimate how peer-to-peer file sharing demand will develop compared to online video and audio (streaming) demand. As we stated in our 2014 report, the churn between legal and illegal media sources depends greatly on legislation as well as availability of legal alternatives. By merging the categories without changing the underlying growth estimates, we are making the assumption that the churn between peer-to-peer and streaming alternatives is a zero sum game with respect to volume demand.

Several studies have since confirmed our conclusion that there will no further uptake of peer to peer file sharing services, and some even find a decrease. Studies by market research firm GfK [15][6] show a 40% penetration of Video on Demand Services among internet users in the 13+ years group in Q4 2015. Sandvine [14] predicts a negative adoption rate in their 2014-2019 forecast for the Western European region. A Norwegian study confirms the strong repressing effect of 'all-you-can-stream' services on banning digital piracy. [9].

3.2.2 Service intensity growth

In our 2014 report, we made some observations on the expected intensity growth of the various services. We argued that an increase in intensity of services leads to and increased demand for speed. With respect to most services, we did not observe strong alterations on our previous estimations. We did make some new observations though

- For *web pages*, growth from intensity can be estimated by looking at the development of the size of websites over time. The HTTP Archive showed that the average total transfer size of websites at the end of 2013 was approximately 1590 KB. At the end of 2015 (two years after our initial publication) the total average transfer size of a web site was about 2190 KB, which indicates a compound annual growth rate of 17.4% [16].
- For *video*, we expected streaming services to deploy a 4K-option in their offerings by 2016, which they now do. Netflix recommends subscriptions with 25 Mbit/s provisioned speed for this type of Ultra-HD quality streams. We argue that the actual bit rate of the video's will be in the range of our earlier predicted 15 Mbit/s, due to compression and buffering.

⁹ See for instance [12].

- For *music*, there does not seem to be a change in the demand for higher-quality audio than currently offered by most streaming services and online audio stores. Observing the highest bitrates offered by popular music streaming services, we see the following exemplary offerings: Spotify - 320 kbps, Google Music - 320 kbps and Apple Music - 256 kbps. This is in line with our earlier claim that there will be little demand for higher than MP3 bitrates. Tidal is the only exception, with the option of lossless 1411 kbps. [17] Their market share however remains very limited.
- For *photos*, we observe newer and smarter services for cloud based storage. Reaching beyond storage, these services are (or will be) able to perform facial and thematic recognition.

3.3 Future revolutionary services

Historical development of internet services has shown that every now and then services appear that were never envisioned. Examples of such services are Google, Gmail, YouTube (introduced in 2005), social networks and, more recently, Netflix.

In 2014, we calculated the expected probability of the development of future services and their expected impact on demand growth. [4] For each service type, we calculated the expected number of times a revolution will occur in the seven-year period. We subsequently modelled the impact of each kind of revolution inversely to the occurrence frequency: high-frequency revolutions have a low impact, and low-frequency revolutions a higher impact. The impact is measured as an increase in yearly traffic in proportion to existing traffic (e.g. an impact of 5% indicates that each year, traffic will grow by an additional 5%). We chose to assign 50% growth and 200% impact to the rarest type of revolutionary event ('once every 50 years') and scaled the impact to the other types using a quadratic and cubic interpolation formula, respectively.

Figure 11 gives a graphical overview of the distribution of the probability of revolutionary events happening and the corresponding impact of such an event on speed demand. By summing the expected growth percentages, it is possible to calculate a compound annual growth rate over the seven-year period. We estimate the year-over-year growth from 'revolutionary' events in the coming seven years will be between 5.5% and 11.1%. The percentages are highly dependent on the estimated impact of the various types of events.

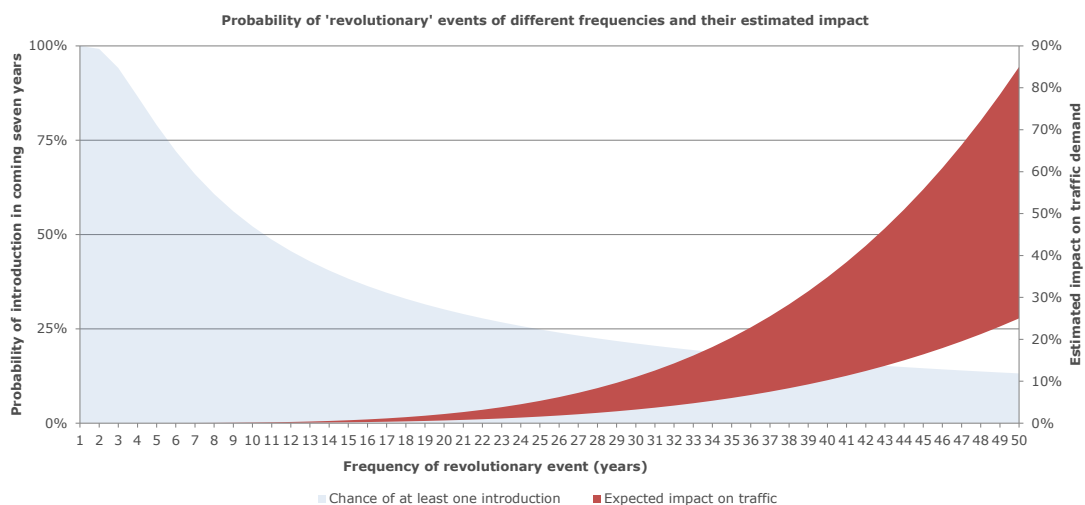


Figure 11. Probability of 'revolutionary' events of different frequencies and their estimated impact

As the time horizon in the 'future services' model is rather large, we do not see reason to change it now, nor do we have new insights that challenge the assumptions made in our earlier study.

3.4 Differences between user groups

Using the different adoption curves for each user group, it is possible to calculate different growth factors per group for each service. Figure 12 gives an overview of the growth rates for the service categories in each of the groups including the growth from increased intensity (assumed to be equal in the groups).

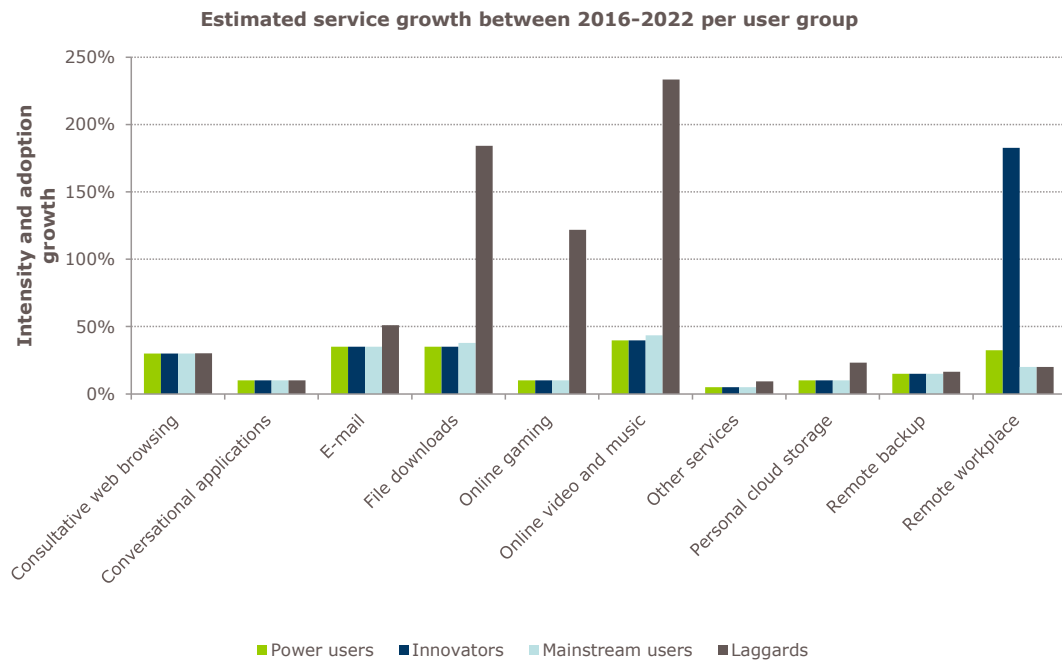


Figure 12. Total growth rate (intensity and adoption growth combined) of services for each user group

Several services are expected to grow significantly according to this figure. Remote backup, conversational applications, online video, online music and personal cloud storage are cases in point. There exist large differences between the user groups. Since the intensity is expected to grow equally among the user groups, the differences arise from the adoption of these services. The graph shows that especially laggards are expected to adopt services on a large scale, in particular remote backup, conversational applications, online video and to a lesser extent online music. Personal cloud storage is on the other hand expected to be adopted by innovators on a large scale. This is because currently the service is assumed to be used mainly by power users and will subsequently dissipate to the innovators group.

4 Future demand

In this chapter we discuss the future demand for traffic and connection speed. As in chapter 2, we start with the aggregate demand for traffic and add further detail by distinguishing service categories. Finally we make the step from traffic volume demand to speed demand by adding a time dimension (urgency/concentration of traffic).

4.1 Aggregate demand

In chapter 3 we project that for the period of 2016-2022, annual growth will be equal to 40.5% for downstream and 44.1% for upstream traffic volume, which equals our estimate in [4] for the period 2013-2020. Figure 13 shows the resulting growth pattern relative to 2016 traffic volumes.

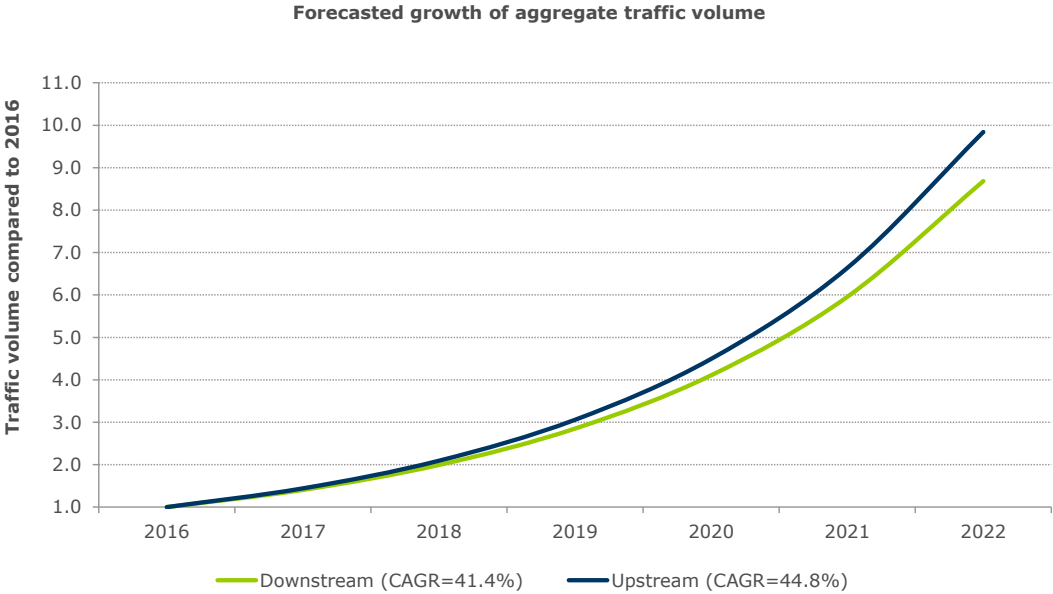


Figure 13. Forecasted growth in aggregate traffic volume demand according to the model. Note that this chart does not imply that there is equal demand for upstream and downstream traffic, as the underlying absolute volumes are different (see Figure 16).

In our model, we make a distinction between growth of existing services, as well as growth following the introduction of future revolutionary services. The growth of existing services alone is responsible for a CAGR of 36.6% (upstream) and 31.6% (downstream) respectively.

4.2 Demand by service

Figure 14 and Figure 15 show the estimated traffic volumes for upload and download respectively generated by existing services, grouped by the different service types.

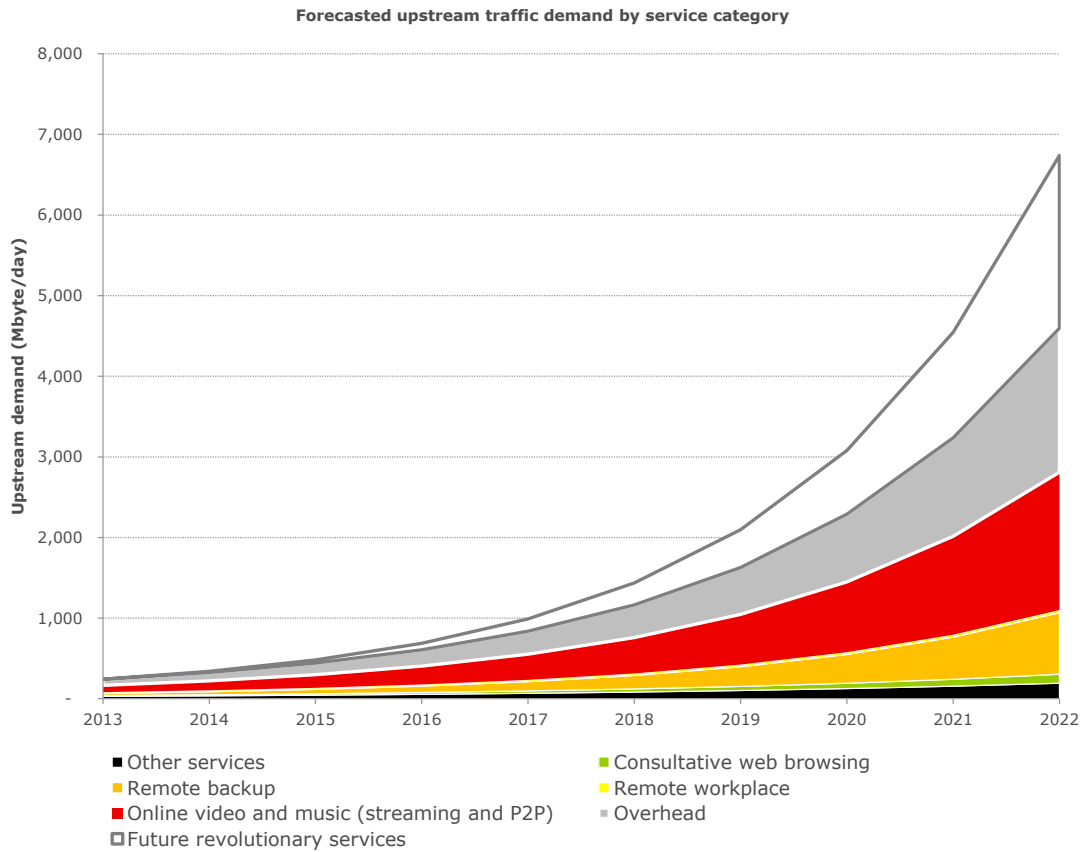


Figure 14. Forecasted average daily upstream traffic volume per residential subscription

Online video and music, online back-up and overhead are the services that will drive upstream traffic volume demand for the coming period of time. As overhead traffic partly consists of acknowledgements of the downstream traffic, this demand is mainly driven by a high download demand. The estimated demand for daily upstream traffic in 2022 will average at just below 7 Gbyte per day.

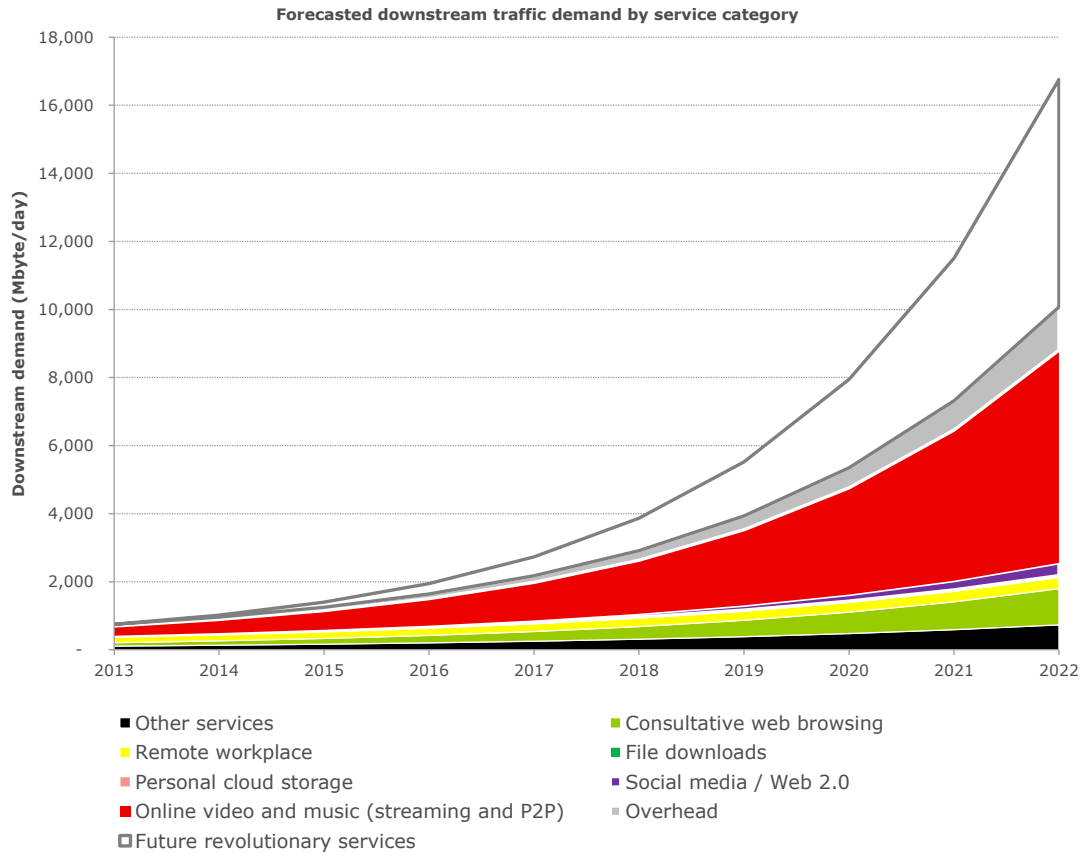


Figure 15. Forecasted average daily downstream traffic volume per residential subscription

Unsurprisingly, online video and music are the major driver of downstream traffic growth. The growth of online video consumption is primarily due to intensity growth (i.e. the move towards HD and higher resolutions) but also growth from adoption by lagging users (primarily driven by the introduction of legal video streaming services such as Netflix).

There is also a major role for future revolutionary services in the downstream direction, which we expect to cause 40% of the total traffic volume demand by 2022. The total downstream demand for 2020 is estimated at 16.6 Gbyte per day, per household.

4.3 Demand for speed

The minimum connection speed required depends on the volume of the traffic to be transferred on a given day, versus the amount of time in which the majority of the transfers takes place. The latter is determined by simultaneous usage of services as well as the duration of service usage. Due to this, ISPs always provision more speed to end-users than the minimum speed required to transfer the daily traffic volume in a day. For example, if a user wants to transfer 200 Mbyte per day, a connection with an average speed of 0.79 Kbit/s would be sufficient to transfer all the traffic within that day. However, assuming that 80% of the 200 Mbytes are transferred within five minutes, the minimum speed required to satisfy the demand is 4.2 Mbit/s. In that case, the ISP may want to provision at least 4.2 Mbit/s. The time in which most (say 80%) of the traffic is transferred is a measure of the 'urgency' of the demand.

We estimate the current urgency of demand by comparing the current daily traffic volume with the currently provisioned connection speeds by ISP A, as well as observed traffic patterns (see paragraph 2.3). Assuming that the urgency of the demand in 2022 will not be different from the urgency in 2016, it is possible to calculate the sufficient provisioned speed in 2022 given the traffic volume predicted.

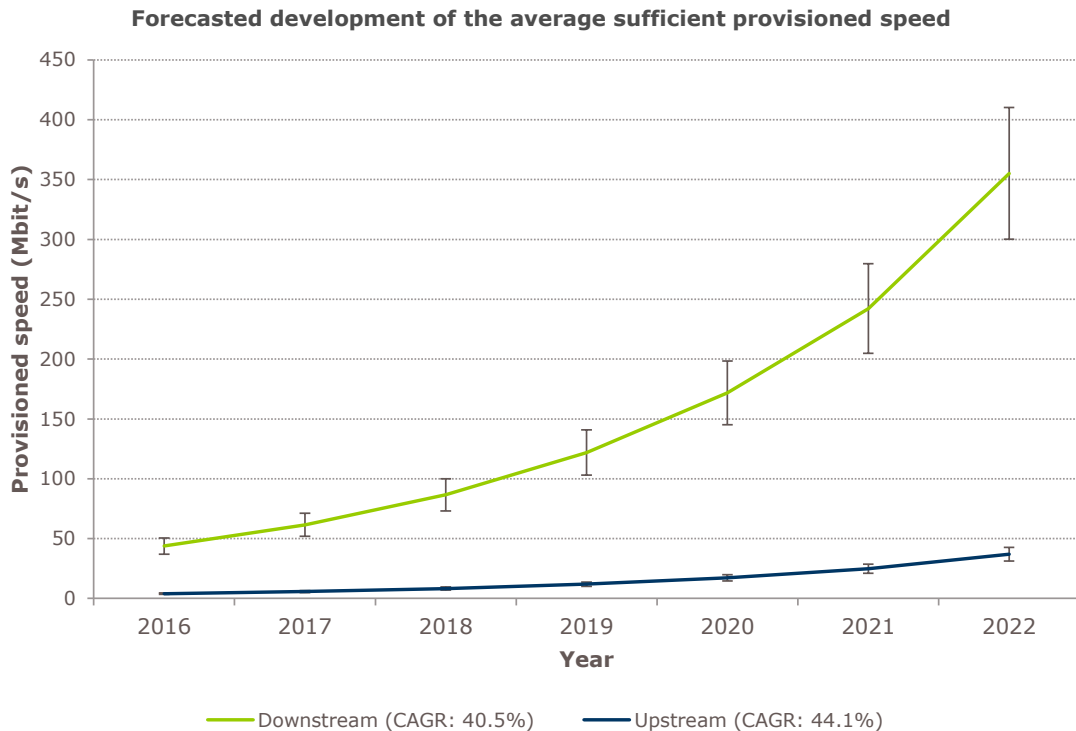


Figure 16. Estimated development of the average sufficient provisioned speed of subscriptions

Figure 16 shows the forecasted development of the average sufficient provisioned subscription speeds using the method described. You can see from the chart that an average subscription will have a sufficient provisioned downstream speed of about 355 Mbit/s in 2020 (compared to 44 Mbit/s in 2016) and an average sufficient provisioned upstream speed of 37 Mbit/s. This estimate is only valid assuming that the current advertised speeds are a reasonable indication of the speed of a 'sufficient' connection. In addition, it is assumed that the urgency of traffic will not change. The error bars in Figure 16 show the speeds required if urgency changes by 20% (i.e. traffic needs to be transferred in 20% more or less time than currently).

4.4 Differences between user groups

4.4.1 Aggregate demand per user group

As the different groups will adopt different services at different rates, the growth pattern is different for each group. The power users have already adopted most services such as personal cloud storage and online back-up but will still see some growth in other areas due to intensity growth. This will create differences between groups regarding the total required upstream versus total required downstream capacity. Figure 17 and Figure 18 show the estimated growth of upstream and downstream respectively for each user group. Note that the y-axis has a logarithmic scale in both figures.

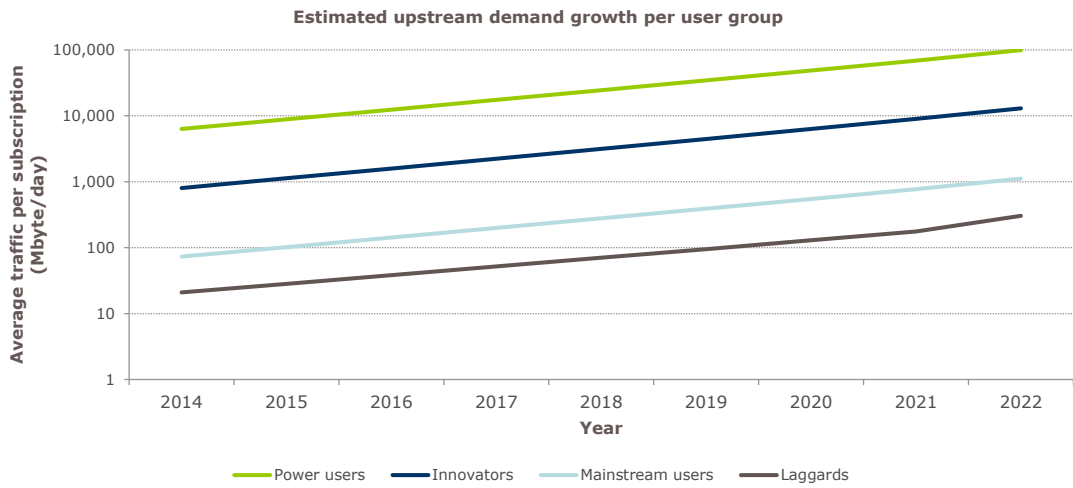


Figure 17. Growth in upstream demand per user group as estimated by the model




Figure 18. Growth in downstream demand per user group as estimated by the model

4.4.2 Demand by service per user group

Table 9 shows the estimated traffic volume demand per subscription in 2022 for the various user groups, broken down into different service categories.

Table 9 Estimated speeds demanded in 2022 for an average subscription in the user group, by service category

	Up Mbyte/day	Down Mbyte/day		Up Mbyte/day	Down Mbyte/day
Power users	93,545	94,012	Innovators	11,430	30,440
Consultative web browsing	2,478	9,459	Consultative web browsing	269	3,036
Conversational applications	351	141	Conversational applications	38	45
E-mail	17	106	E-mail	2	34
File downloads	-	219	File downloads	-	70
Online gaming	1	4	Online gaming	0	1
Online video and music	32,814	30,660	Online video and music	3,565	9,841
Other services	2,726	5,302	Other services	307	1,732
Personal cloud storage	176	169	Personal cloud storage	87	247
Remote backup	8,357	-	Remote backup	908	-
Remote workplace	607	2,430	Remote workplace	66	780
Social media / Web 2.0	1,288	2,533	Social media / Web 2.0	140	813
Overhead	7,543	2,551	Overhead	1,929	830
Future revolutionary services	37,186	40,438	Future revolutionary services	4,119	13,010
Mainstream users	892	4,883	Laggards	275	745
Consultative web browsing	18	508	Consultative web browsing	6	53
Conversational applications	3	8	Conversational applications	3	3
E-mail	0	6	E-mail	0	1
File downloads	-	12	File downloads	-	2
Online gaming	0	0	Online gaming	0	0
Online video and music	247	1,719	Online video and music	19	232
Other services	-	155	Other services	71	153
Personal cloud storage	0	0	Personal cloud storage	0	0
Remote backup	65	-	Remote backup	37	-
Remote workplace	4	130	Remote workplace	1	15
Social media / Web 2.0	10	136	Social media / Web 2.0	4	18
Overhead	285	134	Overhead	55	24
Future revolutionary services	260	2,075	Future revolutionary services	80	246

4.4.3 Capacity and speed demand per user group

Using the earlier described method, we estimated the future sufficient provisioned subscription speeds for each user group. Table 10 gives an overview of these speeds (in Mbit/s).

Table 10. Forecasted average sufficient provisioned speeds (in Mbit/s) for different user groups

		2016	2017	2018	2019	2020	2021	2022
Power users	Up	68	96	135	191	269	379	548
	Down	360	486	656	886	1199	1623	2265
* Note power users:		The estimations for the sufficient provisioned speeds for power users are based on a different method in which traffic for peer-to-peer is modelled to be supply-driven rather than demand-driven. This means that the power users will always maximally utilize the provisioned bandwidth.						
Innovators	Up	9	12	17	25	35	49	71
	Down	111	149	201	271	366	495	693
Mainstream users	Up	1	1	2	2	3	4	6
	Down	17	22	31	42	57	77	109
Laggards	Up	0	0	0	1	1	1	2
	Down	2	3	4	5	7	9	17
All users	Up	4	6	8	12	17	25	37
	Down	44	62	87	122	172	242	355

Appendix A. Methodology

This chapter describes the methodology we applied to answer the research questions presented in chapter 1. We will elaborate on the building blocks used to create the speed demand estimation model.

Overview

In the initial 2014 study, we used three different sources to build the model and estimate its parameters (calibration): measurement of traffic on networks, interviews with experts and literature research. In the update, we relied on repeated measurements of traffic on networks as well as updated literature. In the following paragraphs, we will elaborate on these building blocks and how they are interrelated.

Measurement of traffic on networks

We used (raw) data from measurements of networks that deliver fixed broadband services to residential end users. First of all, we gathered network measurements carried out by several Dutch Internet Service Providers (ISPs), who chose to remain anonymous¹⁰. Throughout this report, we will use the following names to indicate the various ISPs:

- **ISP A:** a large, Dutch ISP providing a variety of primarily cable-based subscriptions. These subscriptions have asymmetric advertised speed limits.
- **ISP B:** a smaller, local Dutch ISP that serves approximately 5,000 subscribers, of whom a small number is a (small or home) business. Connections are either over cable or fibre, but always have symmetric speed limits.

Besides this private ISP data, we used public data sources of the following parties on internet traffic:

- **Sandvine [14]:** Sandvine is a manufacturer of internet traffic monitoring and shaping equipment that presents actual usage data of traffic on its networks every quarter. In the report *Global Internet Phenomena*, they present download and upload traffic in Europe based on actual usage data.
- **Cisco VNI [2][3]:** Hardware manufacturer Cisco publishes the Visual Networking Index online, a tool that estimates IP traffic growth until 2017. It bases its estimates on number of users, application adoption, minutes of use and bitrates and speeds.

Literature research

Literature research provided supplementary insight. We used academic and non-academic sources to provide input for the model estimation, including the following types of literature:

- Several research papers providing detailed insight in the usage pattern of specific services, both quantitative and qualitative.
- Research on innovation studies, adoption models and growth modelling.

¹⁰ Note that the ISPs' data was obtained from network management systems. No data on the content of transmissions was recorded or analyzed, nor could measurements be traced to individual subscribers.

- Additional reports on traffic measurement. Some literature includes types of growth modelling, such as the methodology description of Cisco's VNI [9].

Demand modelling

The first three tasks aimed to provide input for the demand model. The sources were used both to create as well as calibrate the model. The composition of the model, which will be explained in more detail in the following chapter, was based on assumptions extracted from the expert interviews and the literature review. The model was then calibrated with estimates taken from all the building blocks. Significant additional sources were the actual measurements of network traffic by the ISPs, Sandvine and Cisco.

Answering the research questions

Having explained how our speed demand estimation model was created, we will now describe the composition of the model. Firstly, we present a conceptual model, briefly introducing all its components. Later on we will further elaborate on these components and our assumptions. Moreover, we will indicate how the various research sources were used to calibrate the model.

Modelling aggregate demand

The current aggregate demand was estimated as the average amount of traffic transferred by a subscription. In the model, it is assumed that aggregate demand is not determined by supply; that is, the adoption of a particular service is not influenced by the availability of bandwidth. In the Netherlands (and most of Western Europe), experts observe that providers base their decisions on whether to upgrade the network primarily on utilization. Providers continuously monitor their network and desire a certain margin regarding the maximum capacity. They therefore continuously increase the available bandwidth in their attempts to 'stay ahead' of demand. In addition, it is likely that 'power users' who do generate extraordinary amounts of traffic are on the higher-end subscriptions offered by providers. By making this assumption, we can use the measurements obtained for calibration, even though they are about *actual* usage, not demand.

Modelling demand by service

The aggregate demand is a useful concept, but clearly it is comprised of demand for particular services, and these services differ greatly in terms of traffic, speed and capacity. In this study we defined several categories of services, based on precedents from literature as well as expert interviews. In addition to these services, we defined a group of 'other services' which include all services not generally distinguished in measurements, as they are infrequently used or highly specific to particular users. Finally, we defined a category of 'revolutionary services' which consists of all services that cannot be foreseen.

Existing service categories

Based on literature as well as expert interviews, we identified several specific services which are described in more detail in section 4.2.1. To estimate the demand for these concrete services, we used a top-down as well as a bottom-up approach. For some service groups, specific literature was available to estimate this parameter, while for other, no concrete literature sources were found. In that case, we applied a bottom-up approach, meaning that we estimated the traffic for this service based on the traffic needed for one single action, multiplied by estimates for the amount of actions per day and the number of users.

Other services

The category 'other services' consists of services that cannot be measured accurately and are therefore not accounted for in the model. It is not feasible to create a model that captures the demand for each and every service in existence. Thus our aim was to capture at least 80% and to group the remainder, assuming that their influence on future demand is comparable to the average of the other services.

Future revolutionary services

The second special group of services consists of 'future revolutionary services'. These services do not exist yet, but are expected to come into existence and subsequently generate demand. Current demand for such services is, by definition, zero. However, we do expect these services to be developed in the time period analysed, and therefore account for growth in this group.

Overhead

All predictions of traffic apply to the link layer and as such include overhead of the higher network layers. Network measurements from other sources may concern traffic at higher levels and/or include overhead traffic. In the tables and figures in this report, we show traffic for services at the application level and a single item for the overhead, so that the total applies to the link layer.

In the model, two types of overhead are included. The first type of overhead is the overhead to transport application level packets (e.g. TCP/IP headers). This is modelled as a fixed percentage of the application level traffic. The percentage is set to 5%, which follows from averages seen by ISPs. The 5% is also reasonable given the overhead that can be expected theoretically from TCP/IP, the most commonly used protocol that employs acknowledgements.

The second type of overhead is overhead resulting from the need to acknowledge receipt of data. This type of overhead is special and different from the other type of overhead, because it causes traffic in the downstream direction in order to acknowledge packets sent in the upstream direction and vice versa. Depending on the efficiency of the application, this type of overhead traffic can be between 5% to 100% of the traffic affecting the other direction. Following measurements provided by ISP C, we assume that the overhead of this type in the other direction is 10% of the total traffic in that direction, i.e. downloading 100 Mbyte causes 10 Mbyte overhead in the upstream direction.

Modelling differences between user groups

Another useful distinction appeared to be by user group. The experts indicated that there are large differences between types of users. The heavy users generally have a higher traffic demand and in particular much higher upload traffic due to their high use of services that are relatively heavy on upload. Low-end users typically make use of less traffic-heavy services. Consultative web browsing is a case in point: this download-driven service is used by almost every internet user.

For that reason, the model estimates the traffic and speed demand per user group. Moreover, we estimated different adoption figures per user group per service. The next step was to estimate a growth factor for both the usage intensity of services as well as the adoption per user group.

In this study, we defined four categories of users:

- *The power users (2%)*: people who adopt services unusually early and make use of them in ways that are far above the average.
- *The innovators (18%)*: people who are usually early to adopt new services and also make use of most of the features provided by these services.
- *The mainstream (60%)*: the majority group of users.
- *The laggards (20%)*: the group of users that is reluctant to adopt new technologies and services. In general, they use internet services because there is simply no other 'offline' alternative.

Note that the distinction in these groups seems to mirror, at least to some degree, the different types of subscription available from internet service providers. The categories also resemble user categories commonly applied in innovation literature [13].

In order to differentiate growth by adoption between the various groups, a separate adoption curve must be estimated for each group. This is possible because the different groups 'follow' each other in the adoption process: as soon as the power users have adopted a service, the innovators will start adopting, then the mainstream users, and so on. By examining the percentage of adopters required for a group to start adopting (that is, the sum of the sizes of the earlier groups), it is possible to determine the year in which a group will start adopting and the year when this is finished.

From the estimated curve, we see that the power users start adopting as soon as a service is introduced and take four years to fully adopt it (2% of the population has adopted at that point). The innovators start after four years and take six years to adopt (2% + 18% has then adopted). The larger mainstream group adopts in six years as well, after which 80% will have adopted. Finally, sixteen years after introduction, the laggards start adopting and finish in eleven years, and after twenty-seven years, all internet users for which the service is relevant will have adopted. Figure 19 gives a graphical overview of this process by showing the adoption curves for each group. Note that the sum of the separate curves forms a rough approximation of the aggregate adoption curve.

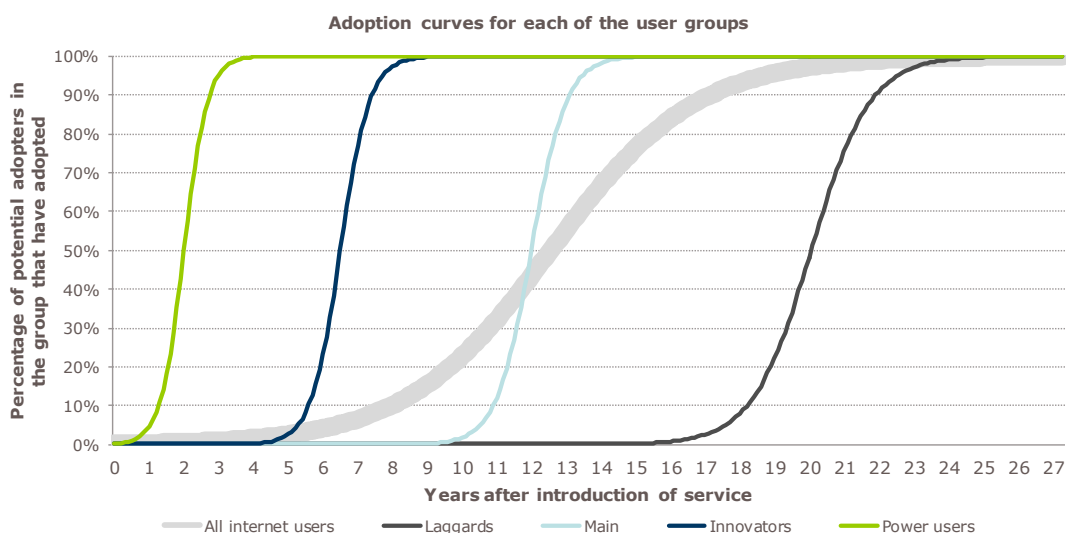


Figure 19. The adoption curves defined for each of the four user groups that add up to the total adoption curve used earlier

Estimating speed demand for a given level of traffic volume

We model the total amount of traffic that will be transferred over an access connection on an average day, which indicates demand for traffic. However, we also aimed to estimate the demand for speed, which is the capacity of the connection in terms of speed (megabits per second) required. For that purpose, we developed a method to translate traffic demand into speed demand. In our model, we assumed that these megabits may be transferred at any time of day. However, users want to be able to use most of the services instantly, not wait every time before their video stream is ready to play or their e-mail is sent. This 'urgency' poses additional requirements for the speed of the connection.

In order to estimate the 'urgency' of the traffic transferred by users, we performed traffic measurements, observing how long users typically take to transfer their daily traffic. We sampled traffic in small time intervals then sorted the samples by size. The amount of samples required to transfer the majority of the traffic (fixed at 80%) gives an indication of the urgency. In the model, we divided the size of the majority of traffic by the time it takes to transfer that amount to obtain the required speed.

Modelling demand growth

Most studies on this topic specify their predictions of various aspects of growth. Although these provide useful data for comparison and cross-validation of our estimates, they are not used as inputs for our model. We limit ourselves to the endogenous growth of fixed, residential connections. The material found during our literature study is however used to perform 'sanity checks' on our estimations.

In order to estimate future demand, we assume that its growth is caused by two different factors. The first is growth by increased adoption: over time, more and more users will get to know and start using existing services. The second is intensity growth of services already used. Users may switch to higher-quality services, or services may start offering higher quality. A good example of intensity growth is a video service starting to offer videos in higher resolutions.¹¹

Future services

In the broadband demand estimation model, it is implicitly assumed that growth will ensue from services that exist today. That is, the resulting growth rates are a sort of 'baseline' that predicts growth, given that no new services will appear outside the currently defined categories. Nevertheless, the history of the internet contains plentiful evidence of new, disruptive services that fitted none of the categories existing up to that point, but have had a huge impact on demand. Examples include services such as Netflix (first introduced in 1999 in the United States) and YouTube (introduced globally in 2005). In addition, 'revolutionary' devices (e.g. the iPhone in 2007 or the iPad in 2010) can cause an increase in traffic demand.

We could model such 'revolutions' simplistically by assuming the introduction of a fixed number of 'revolutions' in a certain timespan, and attaching a certain impact on traffic growth to each of these introductions. However, not all revolutions are equal: every year several services are introduced that were 'unforeseen', but have no significant impact on traffic growth.

¹¹ In these cases, video services often apply a technique that detects a user's bandwidth in real time and adjusts the quality of a video stream accordingly ('adaptive bitrate streaming'). When there is insufficient bandwidth for the high quality stream (e.g. when multiple streams are active at the same time over the same connection), the user's demand is satisfied, but only partially (i.e. with a lower quality stream). In the model, we assume that users always demand the highest quality video stream provided by a service.

Internet banking is one such service, as it was adopted by a large proportion of internet users, but consumes very little traffic. At the other end of the spectrum, there are 'once-in-a-lifetime' revolutions that do have an enormous impact.

By modelling a probability distribution of the impact of revolutions and their expected occurrence frequency, it is possible to find an 'expected impact' of revolutions of different sizes.

Growth by service

In the model, growth by service is demonstrated by two factors: growth by increase in *intensity* and growth by *adoption*.

Intensity refers to the quality received by the end-user. Intensity growth comprises increase in video resolution, sound quality, et cetera. In addition, the user can use a service more intensely than before.

Growth by adoption comes from more and more users starting to use a newly introduced service over time, until all relevant users have been reached. We will now further elaborate on how these two types of growth regarding services were implemented in the model.

Adoption of services

In order to estimate the demand growth caused by service adoption, we model the adoption curves of the services included in the model. The adoption curve describes the proportion of current internet users that start adopting a particular innovation (i.e. using a particular service) over time. The adoption is complete (100%) when all users for which the service is relevant have started using it.

Note that the internet itself has of course not yet been fully adopted by the entire population. While new internet users will contribute to total traffic and speed demand, we are interested in how demand develops at the single subscription level. Therefore, we did not include the adoption of the internet as a whole in our model ('*exogenic*' growth), but rather looked at the adoption of services assuming a user is connected to the internet ('*endogenic*' growth).

In innovation literature, an S-curve has been used successfully to model the adoption of innovations [13]. By collecting data on the current age and adoption among internet users of various internet-based services, we were able to determine the most suitable S-curve. For each service category, we identified the earliest consumer implementation (e.g. globally available to the mainstream) as well as the current adoption. Adoption data were obtained from Eurostat, Cisco and academic publications.

We use the same curve for each service, but vary the position along the curve for each service individually, depending on each service's age and maturity. The year-over-year growth rate of adoption of each service was calculated from the difference in adoption as estimated by the curve during the analysed time period.

Intensity of services

Growth by increased intensity of services is modelled using a single growth factor for each service. We use the expert interviews and literature findings as basis for determining growth factors.

One argument not to be able to extrapolate the trend for higher speed demand is the so-called 'eyes'-argument: given the fact that a household contains a given amount of pairs of eyes (persons), and these pairs of eyes are only able to watch one video stream at a time with a maximum perceivable quality, this implies a maximum needed traffic.

Regarding the growth of intensity of services, a major component is shown to be the growth from technical aspects, such as resolution, colour depth, et cetera. We investigate changes in usage patterns (frequency of use, amount of songs/photos/videos consumed, et cetera) and potential different uses of existing services.

Growth by user group

As described in innovation theory, the adoption of services usually starts with a small group of early adopters, and then gradually spreads through society to reach the mainstream and finally the laggards. As the groups we have defined are based on adoption speed, the growth resulting from adoption should be calculated for each group independently.

The model enables us to estimate what kind of connection will be 'sufficient' to meet the average future demand. Individual users, however, may be further along or less far on the adoption curve, or in general exhibit a higher level of usage intensity. The average connection will be more than enough for one user, but may be limiting for another.

Given the parameters of the groups, we 'break down' our adoption curve for the entire population into distinct curves that model the individual groups. This should of course be done in such a way that the weighted sum of these curves is equal to the original total curve. Splitting the adoption curve can be done by varying several parameters; we chose to vary the maximum adoption percentage and the adoption rate for each group.

Limitations

The model employed in this study estimates future demand for speed by predicting the growth in demand for existing services among different user groups according to the increase in adoption and intensity. In addition, several scenarios for 'revolutionary' services have been included to model the expected introduction of unforeseen services. The advantage of this approach is that it provides a high level of detail as well as several parameters to 'tweak' for example the speed of adoption, magnitude of intensity increase, size of user groups et cetera in order to estimate the impact of various scenarios.

The increase in adoption and intensity have been calibrated using various sources and modelling by the researchers. Different values would lead to a significantly different estimate, and are therefore prone to subjectivity. A case in point is the traffic for peer-to-peer file download. The intensity growth has been calibrated on the historic growth of one ISP, where choosing a different ISP could have resulted in a different estimate.

In the model it is assumed that all traffic generated by a household corresponds to 'actual' or 'intrinsic' demand for a service by an end user. This may not always be the case for two reasons:

1. Some traffic may be 'involuntary'. One example is the traffic generated by viruses and spyware, which a 2013 study has shown to account for up to 61% of traffic of an individual connection [7]. Note that this does not mean that 61% of aggregate total traffic is involuntary, as only a small number of connections will be affected. Several ISPs have put counter-measures in place, and for instance block the connection for an infected subscriber after detecting malware.
2. Demand may sometimes be a result of supply rather than 'actual' demand. For example, peer-to-peer applications generally use all available bandwidth. If the ISP were to raise the speed of their subscriptions, they would see significant increases in traffic volume, which nevertheless do not directly relate to increased *demand*.

When comparing our results with those from other studies, it is important to carefully check how each study deals with these issues.

Finally, a fundamental issue with predictive quantitative models is that it is generally impossible to model the so-called 'black swans', that is, unlikely events that nevertheless can have substantial impact on the outcome. A few examples of such 'black swans' in the context of our study are the following:

- High risk of certain online activities (e.g. illegal file sharing). Users will decide whether to use a certain illegal service based on the 'pay-off' of its usage. Should peer-to-peer file sharing expose users to a high(er) risk from viruses and law enforcement, fewer users will accept the risk.
- Blocking of certain services or content. Several governments have already decided that certain services or websites should not be accessible at all (e.g. file sharing websites) or only with an explicit opt-in (e.g. porn sites). Although such blocking can usually be circumvented by technically skilled users, the majority will simply be unable to gain access.
- Separation or decentralization of the internet. Several (non-free) countries have installed virtual 'Chinese walls' that disallow many types of foreign services. Such separation or decentralization could become a reality in other countries in the coming years following the recent revelations regarding the widespread U.S. intelligence activities.

References

- [1] AMS-IX (2016). *Historical traffic data*. [ams-ix.net]
- [2] Cisco (2011) *Visual networking index (VNI) Global IP Traffic forecast 2010-2015*. Presentation by Mark Nowell. [ieee802.org]
- [3] Cisco (2013). *Cisco Visual Networking Index: Forecast and Methodology. 2012-2017* – white paper
- [4] Dialogic and Eindhoven University of Technology (2014). *Fast forward: how the speed of the internet will develop between now and 2020*. [dialogic.nl]
- [5] EU IPO (20xx). *IP Youth Scoreboard Study*. euipo.europa.eu
- [6] GfK (2015). *Trends in digital entertainment*. [gfk.com]
- [7] Incapsula (2013). Report: bot traffic is up to 61.5% of all website traffic. [incapsula.com]
- [8] Kennisnet (2016). *Handreiking externe connectiviteit*. [kennisnet.nl]
- [9] Music Business Worldwide (2015). *Music Piracy has been 'virtually eliminated' in Norway*. [musicbusinessworldwide.com]
- [10] NationMaster (2014). International Internet Bandwidth in Mbps per capita. [nationmaster.com]
- [11] Pesovic, A. & Sharpe, R. (2012). *Is symmetrical bandwidth a Myth or a Must?* [alcatel-lucent.com]
- [12] Poort et al. (2013) *Baywatch: two Approaches to Measure the Effects of Blocking Access to The Pirate Bay*
- [13] Rogers, E.M. (2005) *Diffusion of Innovations*.
- [14] Sandvine (2013). *The Global Internet Phenomena Report. 2012 H2 – 2015 Q3*.
- [15] Telecompaper (2015). *GfK: een derde gebruikers betaalt voor streaming muziek*. [telecompaper.com]
- [16] The HTTP Archive (2013). *Trends*. [httparchive.org]
- [17] TIDAL (2016). *How good is the sound quality on TIDAL?* [tidal.com]



Contact:

Dialogic
Hooghiemstraplein 33-36
3514 AX Utrecht
The Netherlands
Tel. +31 30 215 05 80
Fax +31 30 215 05 95
www.dialogic.nl/en