



Fast Forward »

How the speed of the internet will
develop between now and 2020

Commissioned by:

NLkabel & Cable Europe

Project:

2013.048

Publication number:

2013.048-1262

Published:

Utrecht, June 2014

Authors:

Ir. Tommy van der Vorst
Ir. ing. Reg. Brennenraedts MBA
Ir. David van Kerkhof
Dr. ir. ing. Rudi Bekkers

Management summary

This study provides insight in the future upload and download bandwidth demand for residential broadband connections, by answering the following research question:

How will upload and download bandwidth change between now and 2020?

Dutch consultancy Dialogic and the Eindhoven University of Technology were commissioned by NLkabel and Cable Europe to carry out this study. Dialogic is an independent research consultancy, focused on innovation and specializing in telecommunications. The Eindhoven University of Technology (TU/e) is a large research and educational institute, and has already collaborated with Dialogic on a variety of projects.

It is important to recognize that residential broadband speed demand has a very diverse character. Households vary considerably in their intensity of use, the type of applications they use their connection for, and the amount of traffic these applications generate. In this study, we address this diversity by differentiating between different categories of users, and our overall outcomes are averages over these categories. While the development of the demand for upload speed is related to that of download speeds, their developments are not necessarily identical. Among other things, download speeds are sometimes more critical to users as they are often 'waiting' to consume content or use an application, while they are typically less pressed when uploading data. However, down- and upload speeds are closely entwined.

In this study, a prospective quantitative model of bandwidth demand was developed in order to answer the research question. Figure 1 shows the estimated rise in demand over the coming seven years, expressed as the average speed of a sufficient subscription. This means that this subscription satisfies the demand of the average user. In 2013, the average sufficient provisioned speeds are estimated as 15.3 Mbit/s (downstream) and around 1.6 Mbit/s (upstream). Note that this is an average: very large differences exist between different user groups. The demand for bandwidth is expected to grow exponentially over the next seven years. The CAGR (compound annual growth rate) is 40% for downstream and 44% for upstream traffic demand. In 2020, sufficient subscription speeds for the average user are forecast to be approximately 165 Mbit/s (downstream) and 20 Mbit/s (upstream).

Forecasted development of the average sufficient provisioned speed

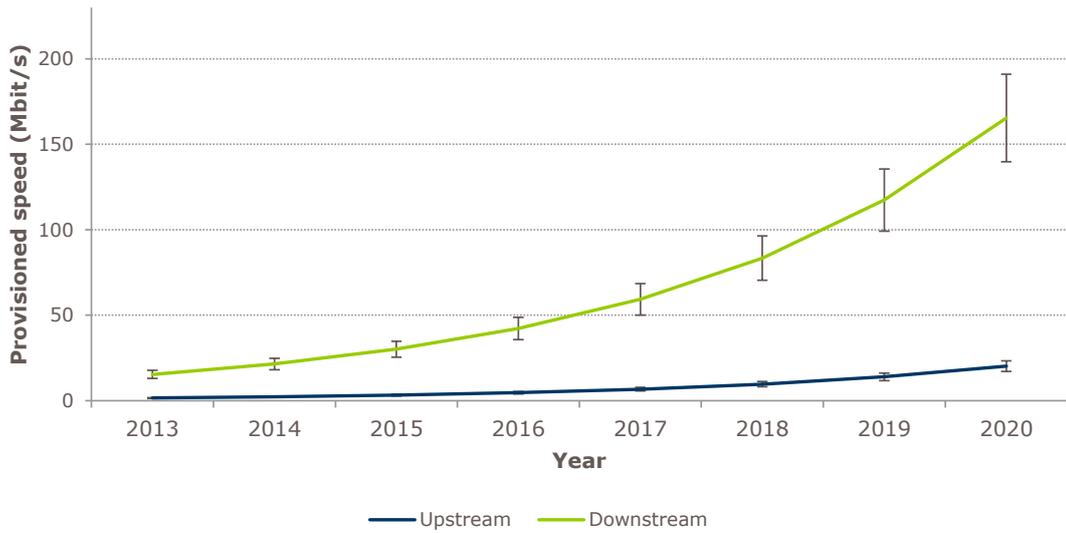


Figure 1. Estimated growth in traffic volume demand relative to 2013 according to our model

The subscription speeds presented in Figure 1 were computed by estimating the total demand for traffic volume per day, and calculating the relative urgency of that traffic (the time in which it is to be transmitted). Traffic volume was estimated per user group and per service category. The following pictures show which services drive the need for more bandwidth.

Estimated upstream demand for services in the period 2013-2020

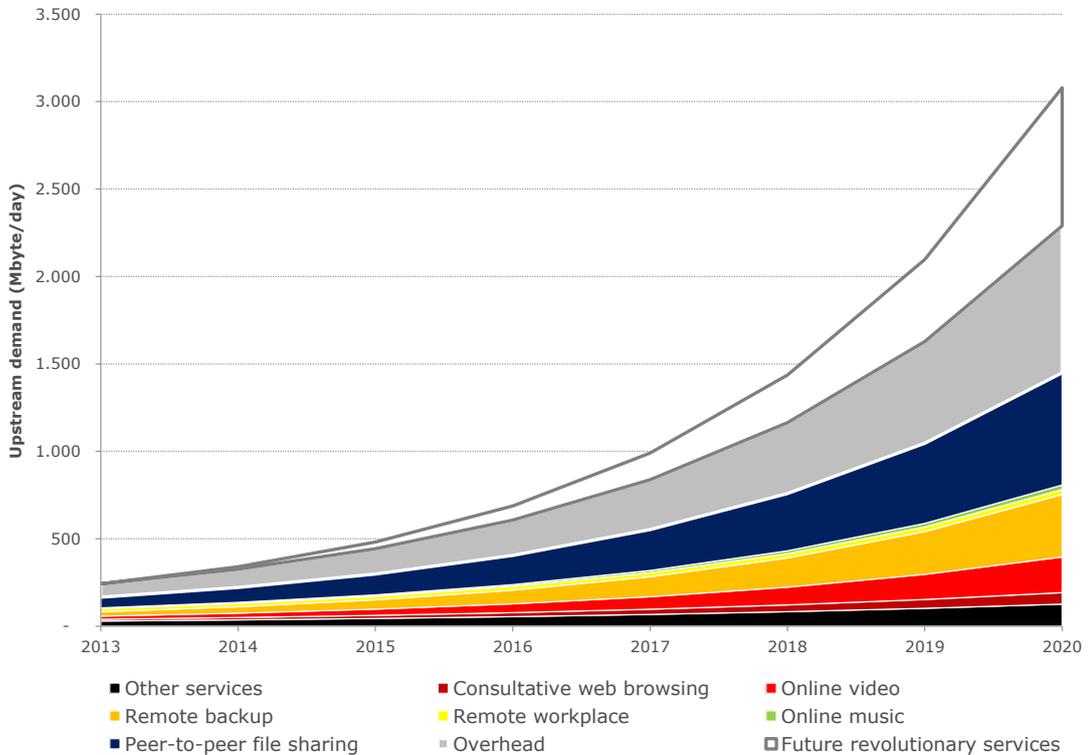


Figure 2. Projected average daily volume of upstream traffic per residential subscription for the years 2013-2020

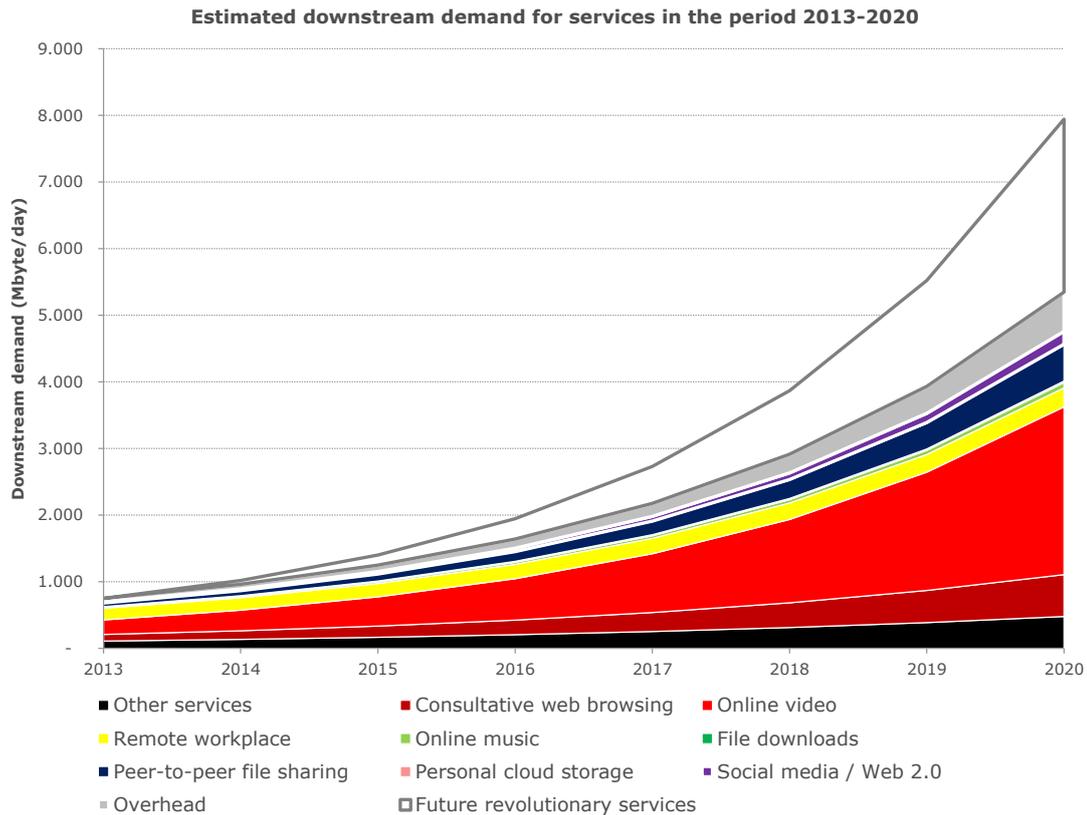


Figure 3. Projected average daily volume of downstream traffic per residential subscription for the years 2013-2020

The first figure above shows the projected volume of upstream traffic demand by service for the period 2013-2020. We estimate that in addition to future revolutionary services, overhead traffic and peer-to-peer file sharing will continue to comprise the majority of the upload traffic. Since overhead traffic partly consists of acknowledgement traffic, this demand is mainly driven by high download demand. Peer-to-peer traffic is modeled to be supply-driven rather than demand-driven, and is therefore subject to the growth of the upstream demand. Other demanding services in 2020 are remote backup services, online video and future revolutionary services. The future revolutionary services are modeled by means of a probability distribution of the impact of revolutions and their expected occurrence frequency. These services could for instance be comprised by a surge in the number of connected devices and accompanying services in a household. The estimated demand for daily upstream traffic in 2020 will average at just over 3,000 Mbyte per day.

Figure 3 presents the estimated volume of downstream traffic demand for the same period. We found that the majority of downstream demand is from online video. Traffic for consultative, overhead and peer-to-peer file sharing are expected to require significant demand. Again, the future revolutionary services also comprise a large part of the traffic in 2020. The total downstream demand for 2020 is estimated at almost 8,000 Mbyte per day.

For our estimates, we distinguished different user groups: power users, innovators, mainstream users and laggards. Take-up of the traffic demand for services is not only caused by the increased intensity of services, but also because more users will start using them.

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

		2013	2014	2015	2016	2017	2018	2019	2020
Power users*	Up	29.2	40.8	57.2	80.2	112.7	158.6	223.4	314.9
	Down	142.4	191.4	257.5	346.9	467.9	631.6	853.5	1154.5
* Note power users:		The estimations for the sufficient provisioned speeds for <u>power users</u> are based on a different method in which traffic for peer-to-peer is modeled to be supply-driven rather than demand-driven. This means that the power users will always maximally utilize the provisioned bandwidth.							
Innovators	Up	3.7	5.2	7.3	10.3	14.4	20.4	28.8	40.8
	Down	44.0	59.0	79.3	106.7	143.7	193.6	261.2	352.7
Mainstream users	Up	0.3	0.5	0.7	0.9	1.3	1.8	2.5	3.5
	Down	6.4	8.7	11.7	15.9	21.6	29.4	40.1	54.6
Laggards	Up	0.1	0.1	0.2	0.2	0.3	0.5	0.6	0.8
	Down	0.8	1.1	1.4	1.9	2.6	3.6	4.9	6.6
All users	Up	1.6	2.2	3.2	4.6	6.7	9.6	13.9	20.1
	Down	15.3	21.4	30.0	42.2	59.3	83.4	117.4	165.4

Table 1 shows the estimated development of subscription speeds that will be sufficient for the various user groups. We found that power users have an extreme usage pattern compared to the other groups. The power users, who account for 2 percent of the users have an estimated demand of 1,155 Mbit/s downstream and 315 Mbit/s upstream in 2020. Laggards exhibit a considerable different usage pattern: they have an estimated downstream demand of 6.6 Mbit/s and an upstream demand of 0.8 Mbit/s in 2020.

A schematic overview of the methodology used in this study is presented in Figure 4. Firstly, we estimated the current aggregate traffic demand and translated this into the current aggregate bandwidth demand. In addition to this translation, we divided the current traffic demand into demand by user group and demand by service. The growth leading to future demand was estimated for each category. We separately estimated the demand for 'revolutionary services' – services that do not yet exist but are expected to in the future.

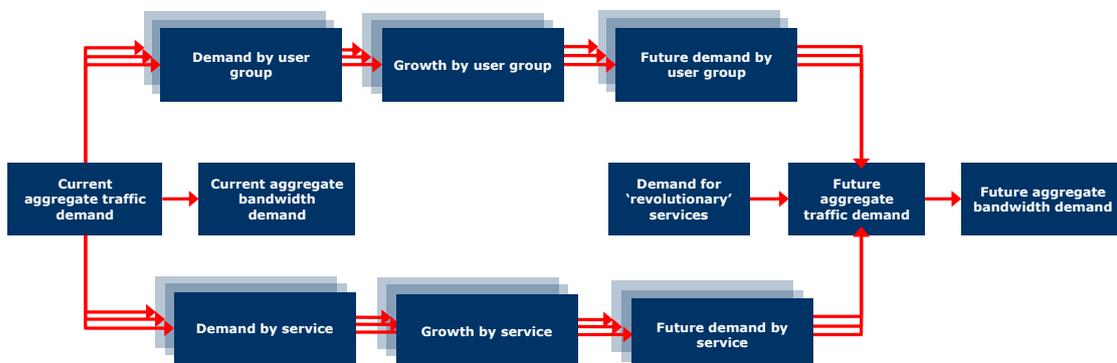


Figure 4. Schematic overview of the demand model

Based on future demand by user group, by service and for revolutionary services, we estimated future aggregate traffic demand. This demand was subsequently translated into

bandwidth demand, analogous to the translation of current traffic demand into bandwidth demand.

The estimates of current demand are based on quantitative data from several sources, mainly Sandvine's Global Internet Phenomena, Cisco's Visual Networking Index (VNI) and three Dutch network operators' recent measurements of cable networks. These measurements were conducted both on Hybrid Fiber-Coaxial networks (HFC) as well as FttH networks. The divisions by service and by user group are based on the literature review and expert interviews, as are the estimates for growth in demand.

Table of contents

Management summary	3
1 Introduction.....	11
1.1 Research questions	11
1.2 Status of the researchers.....	11
1.3 Reading guide	12
2 Methodology	13
2.1 Introduction to methodology	13
2.2 Task 1. Measurement of traffic on networks	13
2.3 Task 2. Interviews with experts	14
2.4 Task 3. Literature research	14
2.5 Task 4. Demand modelling.....	14
2.6 Task 5. Answering the research questions.....	15
3 Model composition	17
3.1 Scope of the model	17
3.2 Modeling approach.....	19
3.3 Modeling aggregate demand	19
3.4 Modeling demand by service	19
3.5 Modeling differences between user groups	20
3.6 Estimating bandwidth demand for traffic.....	21
3.7 Modeling demand growth.....	21
4 Current demand	25
4.1 Aggregate demand.....	25
4.2 Demand by service	26
4.3 Demand for bandwidth	29
4.4 Differences between user groups	31
5 Demand growth	35
5.1 Growth by service.....	35
5.2 Growth in capacity and speed demand	39
5.3 Growth from future revolutionary services	40
5.4 Differences between user groups	41
6 Future demand.....	45
6.1 Aggregate demand.....	45
6.2 Demand by service	45
6.3 Demand for capacity and speed.....	47
6.4 Differences between user groups	49
7 Discussion & conclusions	53
7.1 Answering the research questions.....	53
7.2 Limitations.....	55
7.3 Conclusion	56
Appendix A. Consulted experts	57
References.....	59

1 Introduction

Ever since the introduction of the internet, the demand for bandwidth has continued to grow. Even though home access connections provide generous amounts of bandwidth nowadays, records are still being broken. A recent example is the network traffic record reached when Apple released its latest mobile operating system [13]. The end of growth in bandwidth demand seems to be nowhere in sight.

In this study, we predict the future demand for traffic and bandwidth. We mainly focus on the upstream bandwidth in that respect. However, since the upstream bandwidth is inextricably linked with the downstream bandwidth, the research will also research the latter.

1.1 Research questions

The central question in this study is “*How will upload and download bandwidth demand have developed by 2020?*” 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 upstream traffic?¹
2. To what extent has the need for upstream traffic of currently available applications changed in recent years?
3. Which business applications require high upstream traffic?
4. To what extent will consumers use more business applications by 2020?
5. To what extent will consumers use other applications with a high demand for upstream traffic by 2020?
6. Which upstream and downstream speeds will be sufficient for future demand?

Our research focuses mainly on consumers in the Netherlands and other West European countries with highly developed broadband markets. Questions 1, 2 and 3 focus on the current situation, while the final three questions deal with the future development of consumer demand for internet traffic.

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 (until 2020).

1.2 Status of the researchers

This study was conducted by Dialogic *innovatie & interactie* and the Eindhoven University of Technology (TU/e). Dialogic is a research consultancy in the Netherlands, focused on innovation and specializing 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. Examples are ACM (the Dutch national regulatory

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

authority for telecommunications), European Commission, Dutch Ministry of Economic Affairs, most Dutch provinces, et cetera. Dialogic maintains close connections with the world-wide academic community, and has already collaborated with Eindhoven University of Technology on various telecommunications-related projects. The Eindhoven University of Technology (TU/e) is a large research and educational institute based in the south of the Netherlands, and has already collaborated with Dialogic on a variety of projects.

1.3 Reading guide

The report is structured as follows: Chapter 2 outlines our methodology. In order to find answers to the research questions, we conducted interviews with experts, desk research/literature study and used measurement data on upstream and downstream traffic. In chapter 3 we elaborate on the composition of the bandwidth demand estimation model. In the subsequent three chapters, we look at current demand, estimating firstly the growth factors involved and finally future demand. Chapter 7 ends with our conclusions, discussion and limitations.

2 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 bandwidth demand estimation model.

2.1 Introduction to methodology

As you can see from the diagram in Figure 5, 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.

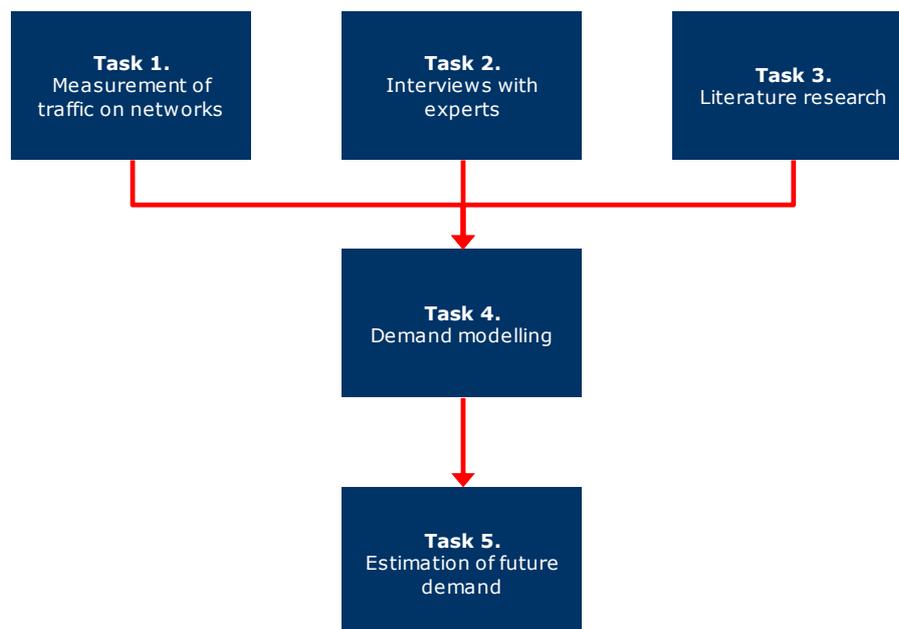


Figure 5. Overview of the research methodology

In the following paragraphs, we will elaborate on these building blocks and how they are interrelated.

2.2 Task 1. 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 anonymous². 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. Traffic measurements were conducted on random days from November 5-12 in 2013 on a randomly selected subset of 1000 subscribers.

² 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.

- **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 fiber, but always have symmetric speed limits. Measurements were made on November 23, 2013 and cover all subscriptions.
- **ISP C:** a large, Dutch ISP providing a variety of primarily cable-based subscriptions. These have asymmetric advertised speed limits.

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

- **Sandvine [2][3]:** 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 [7]:** 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.

2.3 Task 2. Interviews with experts

We conducted ten interviews with international domain experts to confirm and validate expected trends regarding demand and usage, for both upload and download speeds and traffic. The ten experts fall into three groups: service providers, researchers / academics and network operators. Interviewees were selected by the project team based on their expected expertise and represent a balanced combination of the above-mentioned expert types. Half of the interviewees are Dutch; the other experts are mostly West European, and one American. In some cases a relevant company was approached and the company selected the most suitable staff member; in other cases the project team deemed a person suitable, based on published research or the project team's previous experience. A list of the interviewees can be found in Appendix A. The interviews were semi-structured and lasted on average 1.5 hours.

2.4 Task 3. 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.
- Additional reports on traffic measurement. Some literature includes types of growth modelling, such as the methodology description of Cisco's VNI [9].

2.5 Task 4. 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.

2.6 Task 5. Answering the research questions

Finally, the outcomes of the demand model were used to answer the research questions presented in chapter 1.

3 Model composition

Having explained how our bandwidth 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 (the calibration is presented in chapter 5).

3.1 Scope of the model

For the purposes of this study, we made several key choices regarding the model's scope and boundaries. The aim was to estimate bandwidth demand by the year 2020 and the development between 2013 and then. The unit of analysis was household connections. Although we did not model individual households, we aimed 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; also small offices or home offices (SOHO) are excluded.

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

We have modelled demand for bandwidth rather than availability of bandwidth, i.e. bandwidth traffic consumption is driven by demand rather than supply. We assumed 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.

3.1.1 Traffic and bandwidth

In this report, we frequently refer to the terms 'traffic' and 'bandwidth'. 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 *bandwidth* of a connection refers to the amount of traffic it can transfer in a given time period (usually measured in megabits per second).

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 bandwidth 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 bandwidth 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 6 illustrates the relationship between bandwidth and traffic by showing the bandwidth required to transfer various amounts of data in various time periods.

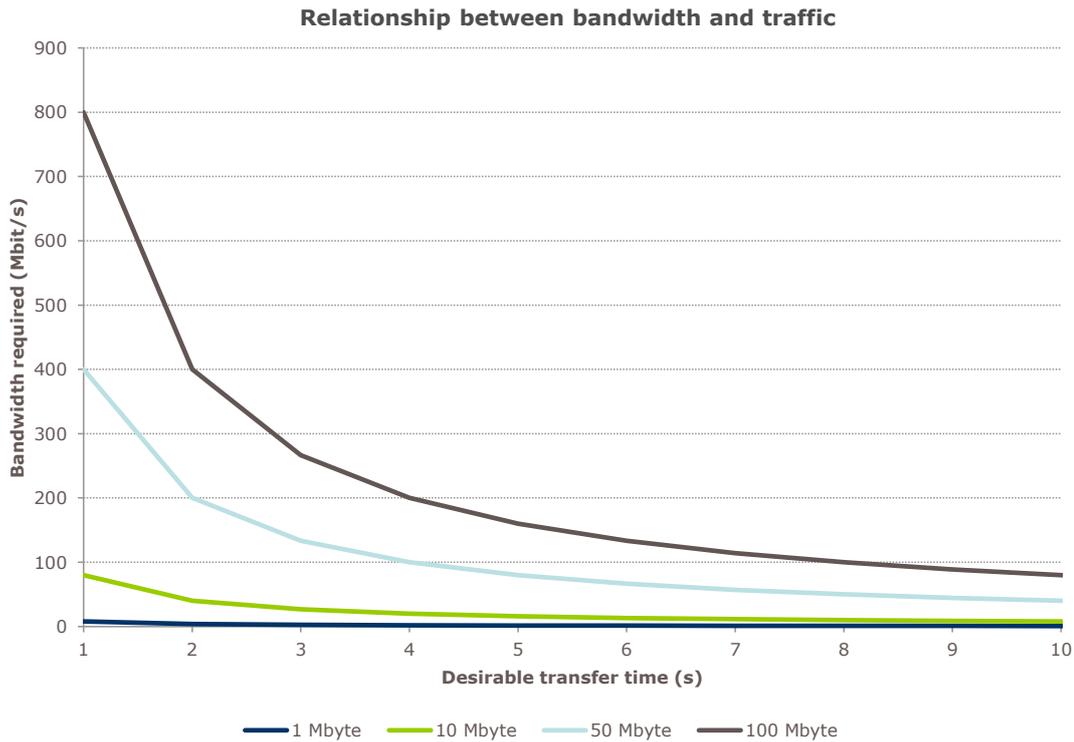


Figure 6. The relationship between bandwidth and traffic

Here, bandwidth is always expressed in megabits per second (Mbit/s) whereas traffic is expressed in megabytes per day or month (Mbyte/day, Mbyte/month). 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 modeled 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.

3.2 Modeling approach

The bandwidth 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 7).

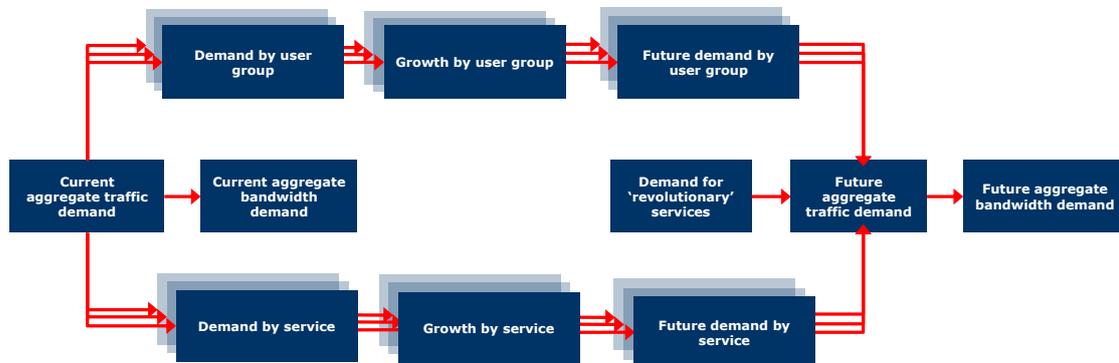


Figure 7. Schematic overview of the demand model

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 bandwidth demand, the final step was to estimate bandwidth demand based on current traffic demand. This was done for both current and future demand.

In the next paragraphs, we will elaborate on the various steps to create the demand estimation model.

3.3 Modeling 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.

3.4 Modeling 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. Paragraph 4.2.1 lists service categories based on currently

available services. 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.

3.4.1 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.

3.4.2 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.

3.4.3 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 analyzed, and therefore account for growth in this group.

3.5 Modeling 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 bandwidth 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 [10].

3.6 Estimating bandwidth demand for traffic

We modeled 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 bandwidth, 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 bandwidth 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 bandwidth 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 bandwidth.

3.7 Modeling demand growth

Earlier literature and the experts agree that the demand for traffic and bandwidth will grow in the coming years. Whether this will continue until the end of time is contested by the experts, but all agree that growth until 2020 is indisputable. Several earlier studies estimate a year-over-year growth factor of bandwidth demand. Cases in point are Nielsen's Law of Internet Bandwidth [12], which states that a high-end user's connection speed grows by 50% per year, and more recently Dialogic and TNO (Netherlands Organisation for Applied Scientific Research) have estimated 40% year-over-year growth [8].

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 have not been used as inputs for our model. We limit ourselves to the endogenous growth of fixed, residential connections. The material found during our literature study was however used to perform 'sanity checks' on our estimations.

In order to estimate future demand, we assumed 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.³

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.

From our interviews with experts, we compiled the following (non-exhaustive) list of services and events that could fit the above description of 'revolutionary':

- A surge in the number of connected devices and accompanying services in a household
- The introduction of advanced, thorough surveillance applications, where all human actions are monitored centrally
- The widespread adoption of domotics technologies. A key component is video surveillance
- A large take-up of distributed computing, in which personal devices solve computational problems (comparable to Seti@Home)
- The widespread adoption of a digital currency (e.g. Bitcoin)

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. 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 modeling 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. In paragraph 5.3 we explain how we operationalized this method.

³ 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.

3.7.1 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 modeled 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 bandwidth 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 [10]. 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 used the same curve for each service, but varied 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 analyzed time period.

Intensity of services

Growth by increased intensity of services was modeled using a single growth factor for each service. We used the expert interviews and literature as basis for determining growth factors. The interviews with experts highlighted several global trends regarding the increase in intensity of services. The main conclusion is that all experts expect a continuous growth of the amount of traffic for the coming years. They acknowledged a certain amount of uncertainty according to this prediction, but were confident enough to make this claim for the time period modeled.

The interviews revolved around the question whether there could be a maximum demand. One argument not to be able to extrapolate the trend for higher bandwidth 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. However, the experts agree that for the coming five years, this maximum will not be reached: first of all the number of parallel video streams is still growing (also see the growth in adoption), and moreover the video quality is still expected to improve, for

instance the increased resolution of video in terms of 4K or 8K High Definition Television is becoming more widespread.⁴ This discussion even neglects the uptake of revolutionary services. A case in point is the increase of connected devices per household that only engage in machine-to-machine interaction.

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

The experts also mentioned specific services that they do not expect to grow in intensity. Music streaming is one such case: the current quality is already so high that an even higher quality will not result in a better listening experience. Voice traffic is another example that is not expected to increase in terms of intensity.

3.7.2 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 enabled 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 were able to '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.

From our interviews with experts and statistics from Eurostat and the Dutch Statistics Institute, we identified several trends and figures on adoption. We found for instance consultative web browsing and e-mail to be used by almost all internet users. The interviewed experts expect online video and music streaming to become more and more popular, sometimes even substituting the linear TV product.

Another reported trend is that an increase in the usability of services initially only used by power users, will make them available to the technologically less savvy users. A prime example is cloud storage: thanks to easy-to-use services such as Apple Photo Stream – which automatically synchronizes photographs taken on mobile devices with iCloud to be accessed from other devices – less knowledgeable users are making use of these services.

⁴ Also compare to the planning in the mobile domain: the Electronic Communication Committee [17] indicates multimedia is expected to account for a rising proportion of traffic in the future, as the proportion of video traffic in networks grows.

4 Current demand

We now turn to modelling the current demand for bandwidth. Starting with the aggregate demand (which concerns all internet users), we provide further details by distinguishing the various services. By adding a time dimension, we could then make the step from traffic volume demand to bandwidth demand. This enabled us to finally estimate the current bandwidth demand for the various user groups.

4.1 Aggregate demand

Table 2 is an overview of the measurements of aggregate total traffic we collected. The aggregate traffic is specified as the average per subscription, as well as the median traffic for subscriptions. The difference between these measurements indicates the skewness of distribution of traffic over subscriptions. If the median is (much) lower than the average, a minority of subscriptions is responsible for a large share of total traffic.

Table 2. Average total monthly volume of traffic per household according to various sources (n/a denotes data was not available; figures denoted with '+' were specified in gigabytes and therefore less precise)

Measurement			Monthly traffic per subscription (Mbyte)			
ISP	Period	Location	Up		Down	
			Average	Median	Average	Median
Sandvine	2013 H1	Western Europe	2,500+	486	10,900+	5,300+
Sandvine	2013 H2	Western Europe	3,400+	751	14,000+	6,300+
ISP A	2013-11	The Netherlands	7,466	622	n/a	n/a
ISP B	2013-11	The Netherlands	2,655	n/a	6,042	n/a

According to Sandvine (manufacturer of internet traffic monitoring and shaping equipment), for the second half of 2013 the total monthly consumption of bandwidth averaged at 3.4 GB upstream and 14.0 GB downstream per fixed, residential connection in Western Europe [3]. In the first half of 2013, Sandvine had reported an average consumption of 2.5 GB upstream and 10.9 GB downstream [2]. This indicates a growth of 36% and 28% in six months for upstream and downstream bandwidth, respectively, which is substantial.

In the Netherlands, bandwidth consumption seems to be above the West-European average [4]. Statistics however vary greatly. ISP A figures indicate that the current average monthly total upstream traffic generated by a residential connection is approximately 7.5 GB. Other data from ISP B show an average of about 6GB total monthly downstream and 2.6 GB total monthly upstream traffic, which seems more in line with Sandvine's figures. The significant difference between the Dutch ISPs' averages is curious, as these cover the same time period and a comparable population. Possible explanations for the differences are:

- Both measurements are based on a relatively small sample (about a thousand randomly selected subscribers in one case and all five thousand subscribers in another). In small samples with high skewness, the average is unstable. Removing a

single user from the ISP A data set lowers the average by about 13%. The median is a more stable measurement given the skewness of the distribution, but was unfortunately not available from all sources.

- There may be behavioral differences between the ISP A and ISP B subscribers because these ISPs cover different geographical areas (one has perhaps more urban subscribers than the other). Such differences can affect the adoption of services as well as the intensity of use.
- The measurements were conducted by the ISPs themselves in various ways, and the variances may have been caused by technical differences (e.g. the inclusion or exclusion of certain overhead traffic). IPTV traffic is excluded in both measurements, but certain types of on-demand video services may have been included.
- The Sandvine measurements include DSL connections, which are on average slower than the types of connections (cable, fiber) provided by the ISPs from which traffic measurements were used in this study (see Table 4).

4.2 Demand by service

Online traffic is generated by a multitude of different internet-based services. Some of these services are more than twenty years old (e.g. e-mail or the WWW) whereas other services have only just been introduced. As the growth potential for one service differs from the other, in order to reasonably predict traffic growth, it is therefore necessary to distinguish these services. Business use is not captured as a separate category, but comprised in the following categories defined below: Conversational applications, remote workplace and work file access, and personal cloud storage and file synchronization are types of services also used for business purposes.

4.2.1 Service categories

From literature and interviews with experts, we have identified several service categories:

- **Consultative web browsing:** all activities related to obtaining information. Usually, this is done through the World Wide Web (WWW).
- **E-mail:** sending and receiving e-mail messages. Note that web-based e-mail is not included in most e-mail traffic measurements, as it is instead placed in the 'Web 2.0' category.
- **Social media and 'Web 2.0':** all online services where sharing user-generated content with other users plays a central role. This category includes social networks like Facebook and Twitter, as well as review websites and market places (eBay and AirBnB).
- **Remote backup:** services that allow consumers to periodically make a back-up of their files (in some cases restricted to certain types of files, e.g. photos or documents), only meant to be retrieved in an emergency. Examples are Carbonite and Backblaze.
- **Conversational applications:** services that provide text conversation, voice and/or video calls between two or more parties. Current popular examples include Skype, WhatsApp, Apple FaceTime and Google Hangouts. These types of applications are also used for business use.

- **Online video:** any service that provides non-linear video streaming (users can start watching what they want at any time). The best known services are YouTube and (more recently) Netflix.
- **Remote workplace and work file access:** secured access to documents and other services normally only accessible in the workplace. This is a typical business application.
- **Online music:** any form of online music streaming or download service. This includes online music shops (Apple iTunes Store) as well as streaming services (Spotify, Deezer).
- **File download:** any type of download not already included in other categories. Examples are app/software downloads and updates and OS downloads and updates such as Windows Update and Steam.
- **Online gaming:** service that allows the end-user to play games against or with other end-users through the internet. An example is the Blizzard server, facilitating games like World of Warcraft, Diablo and Starcraft.
- **Peer-to-peer file sharing:** service that allows files to be shared between end-users from their home computer(s). Older examples are Napster or Kazaa and a contemporary one is BitTorrent.
- **Personal cloud storage and file synchronization:** a service that allows files to be stored at a centralized location, from where it is accessible from a plurality of device types, only requiring an internet connection and certain software. An example is Dropbox. This service could also be used for business applications.
- **Other services:** any service that does not fall into one of the categories defined above. An example could be the offloading of mobile traffic by means of a femtocell or WiFi.
- **Future revolutionary services:** any service that does not yet exist, but may come into existence, and generate demand, in the time period analyzed. Current demand for future services is set at zero (as there is no reason for users to require bandwidth for services that do not exist).

4.2.2 Estimating the demand per service

Various sources were used to accurately model the distribution of traffic demand over services. We distinguished two broad strategies: top-down sources (where traffic is measured and subsequently split into different categories) and bottom-up sources (in which the bandwidth usage of separate services is calculated and subsequently summed). The next two paragraphs explain these sources in more detail.

Top-down traffic statistics

The Sandvine 2013 'Global internet phenomena' reports [2][3] contain aggregate traffic figures on residential usage in Western Europe. The Cisco VNI reports provide several measurements on online gaming and video traffic. Finally, a study commissioned by the Broadband Stakeholder Group (BSG) conducted a similar exercise to forecast bandwidth demand [4]. This study focused somewhat more on audio and video services, and used a

different forecasting approach. In addition to the literature, we used data from traffic measurements performed on ISP A's network.

For the first half of 2013, Sandvine attributes 12.7% of upstream traffic to consultative web browsing, 4.7% to social networking, 8.1% to communications and 51.4% to peer-to-peer applications. For downstream traffic, 27.8% is attributed to web browsing, 5.3% to social networking, 2.9% to file downloads and 14.9% to peer-to-peer downloads [2]. These figures refer to usage at peak periods; it is assumed that the distribution of traffic does not differ greatly in the less busy periods; despite reasons to suspect this is not the case for certain services, we decided to correct for this issue at a later stage. We converted the percentages provided by Sandvine to absolute traffic volumes by using the average volume totals in their report.

Figure 8 shows the distribution of upstream traffic as measured by ISP A and Sandvine. A difficulty with comparing the measurements is that the definitions of categories differ slightly between sources (e.g. Sandvine attributes 4.76% to 'social networking', which we had to add to the 'browsing' category to make the figures comparable with ISP A's data), and therefore do not correspond with the service categories defined above.

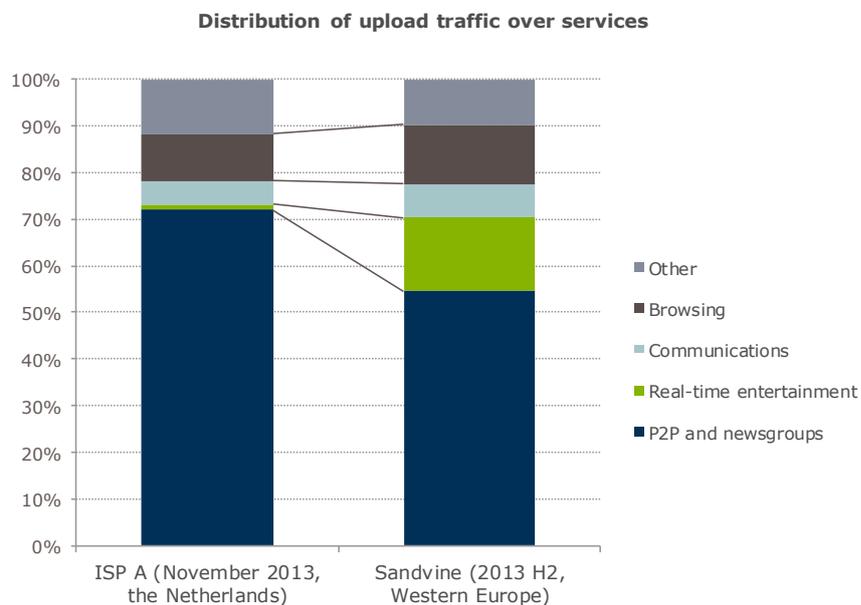


Figure 8. Distribution of upstream traffic in the second half of 2013 as measured by ISP A and Sandvine

Note that more than 70% of the upload traffic of ISP A is comprised by traffic from peer-to-peer traffic. This appears to be a typical pattern for the usage of upload traffic: ISPs report that regardless of the upstream capacity, a small group of extreme users always fills their capacity. For that reason, this usage does not seem to reflect demand, but is rather supply-driven. Note that the traffic for newsgroups is negligible in the upstream: newsgroups primarily impacts downstream traffic.

Bottom-up approximations

In addition to traffic measurements, we also used approximations based on technical aspects and requirements of the various services. For example, the amount of bandwidth used for remote workplace and work file access was approximated using data from

Microsoft on (a) the bandwidth requirements for remote desktop technology [6] and (b) assuming that on half of the workdays there will be a single person working eight hours on a remote desktop. The amount of traffic to be transferred for incremental online back-up was based on the assumption that the average household produces about 10 megabytes of information that should be backed-up (of which photos are probably the heaviest files). Note that the calculated amounts were also corrected for adoption (not all users use these services).

4.2.3 Modeling demand by service

For the final version of the model, we compared our bottom-up approximation with the top-down figures reported by Sandvine and the ISPs. We combined measurements by taking their (weighted) average. Where the figures differed significantly, we instead approximated the traffic bottom-up as described above. We subsequently chose the top-down measurement that was closest to our approximation.

For the 'other services' category, we averaged the residuals from the different sources. In the model, we assumed 12% of downstream traffic and 10% of upstream traffic is unaccounted for by the other categories.

Table 3 shows the average daily total traffic generated in 2013 per subscription as used in the model, divided by these service categories. Note that the totals in the models are very close to the aggregate averages provided by Sandvine [3].

Table 3. Services and the average amount of traffic generated per subscription in 2013 as used in the model

Service	Average daily volume (Mbyte)	
	Up	Down
Consultative web browsing	10.4	99.4
E-mail	0.3	5.0
Social media / Web 2.0	3.9	18.9
Remote backup	25.0	0.0
Conversational applications	6.6	6.6
Online video	17.6	217.5
Remote workplace	17.4	174.5
Online music	5.0	15.0
File downloads	0.0	10.4
Online gaming	0.0	0.1
Peer-to-peer file sharing	60.9	51.1
Personal cloud storage	1.2	3.0
<i>Other services</i>	16.5	82.0
<i>Overhead</i>	76.6	50.7
Net total	241.5	734.2
Net total (monthly)	7,365.8	22,393.2
Net total (as measured by Sandvine in 2013 H2 for Western Europe)	3,481.0	14,336.0

4.3 Demand for bandwidth

Figure 9 shows the average amount of internet traffic transferred by consumers, measured in hourly intervals on a random day in November 2013 on the ISP B network. We can

clearly see that most users are using online services in daytime, which is not surprising. Downstream traffic increases gradually over the course of a day and peaks at around 7:30 pm. A similar pattern can be 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).

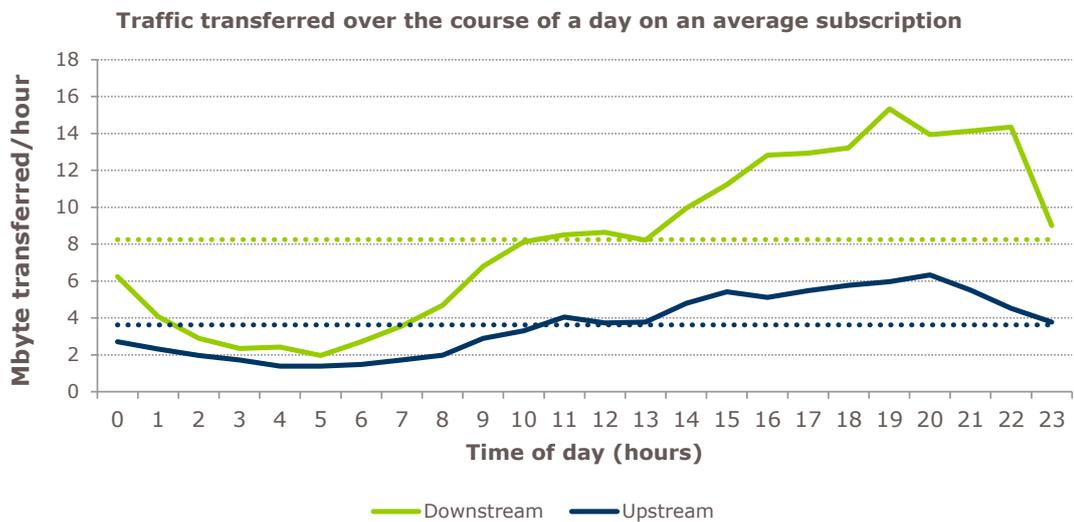


Figure 9. Average volume of traffic transferred on an access connection per hour (measured on the ISP B network in November 2013).

Figure 10 shows how the total daily traffic volume per user is concentrated in a particular time of day. The chart is based on the same data as in Figure 9. According to the data, 80% of the download traffic is, on average, consumed in 819 minutes of a day, which amounts to around 57% of a day. The 80% majority of upstream traffic is consumed in 879 minutes, about 61% of a day.

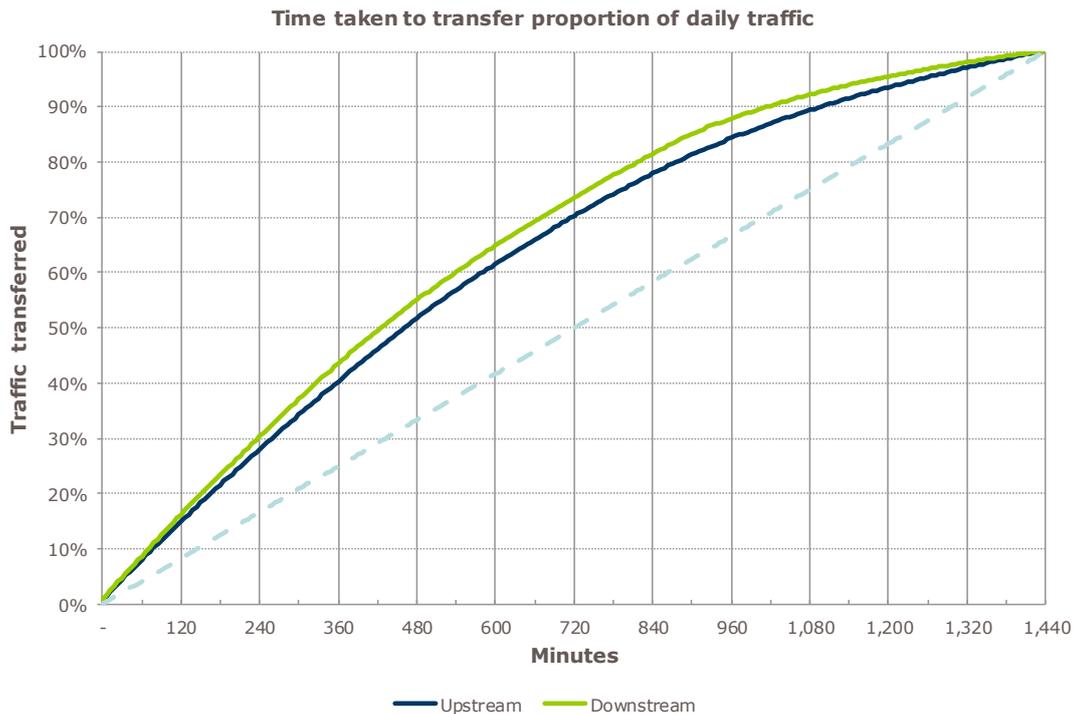


Figure 10. The cumulative proportion of total traffic volume transferred in a day by minutes (measured on the ISP B network in November 2013)

We can see from the above graph that on average, consumers take 60 minutes ($1/24^{\text{th}}$ of a day) longer to upload 80% of the daily traffic they produce than to download the daily amount of traffic consumed. In other words, the bandwidth required for upstream traffic more closely resembles the bandwidth required when the consumption of traffic is equally distributed over the day (depicted with the dotted line in Figure 10) than the bandwidth required for downstream traffic. This leads to a difference in the minimum connection speed required between upstream and downstream.

There are several explanations for the pattern observed above. The most likely one is that many services that are among the top upstream bandwidth users (e.g. Dropbox, online back-up, peer-to-peer file sharing) primarily work in the background, whereas download-intensive services are 'foreground' services for which the end user is actually waiting. An end user does not usually stay on his computer to watch a peer-to-peer program finish an upload to other users, or wait for his Dropbox to finally catch up with the latest version of all the files on the server (although the user could of course be interested in particular files). A downstream-intensive service, on the other hand, usually generates traffic at the initiative of the end user: a video stream is started or a file is being actively downloaded.

4.4 Differences between user groups

Table 4 shows the various ISP subscriptions, the associated speed limits and subscription numbers (where we were allowed to share this information). Note that the actual speeds

which can be attained are influenced not only by speed limitations, but also by other factors.⁵

Table 4. Information on subscriptions from the ISPs whose traffic measurements were used (n/a indicates no data was available).

Subscription	Number of subscriptions	Advertised speed (Mbit/s)	
		Up	Down
ISP A			
Medium	n/a	1-5	10-50
High	n/a	1-5	50-100
Extreme	n/a	5-10	100+
ISP B			
Basic	14%	1-5	1-5
Low	7%	5-10	5-10
Medium	72%	10-50	10-50
High	7%	50-100	50-100
Extreme	1%	100+	100+
ISP C			
Medium	40%	1-5	10-50
High	53%	5-10	50-100
Extreme	7%	10-50	100+

Figure 11 gives an overview of the distribution of upstream traffic over the subscribers of ISP A. The pattern in the figure was also found in the other ISPs' data.⁶ The chart shows the average over several days, as usage is expected to differ between weekdays and weekends. The distribution for ISP A, although closely resembling the distribution in Sandvine's 2013 H2 report, is a bit more skewed (80% of traffic is consumed by about 8% of subscribers to ISP A, whereas Sandvine states it is consumed by around 12% of subscribers). Interestingly, the differences between different days are rather small and statistically insignificant given the same size.

⁵ One of the factors that has the greatest impact on connection performance is overbooking. Regardless of technology, overbooking is common practice in consumer access networks as it is necessary to make the networks cost-effective.

⁶ In this report we decided to show the pattern for ISP A because it had the most detailed data.

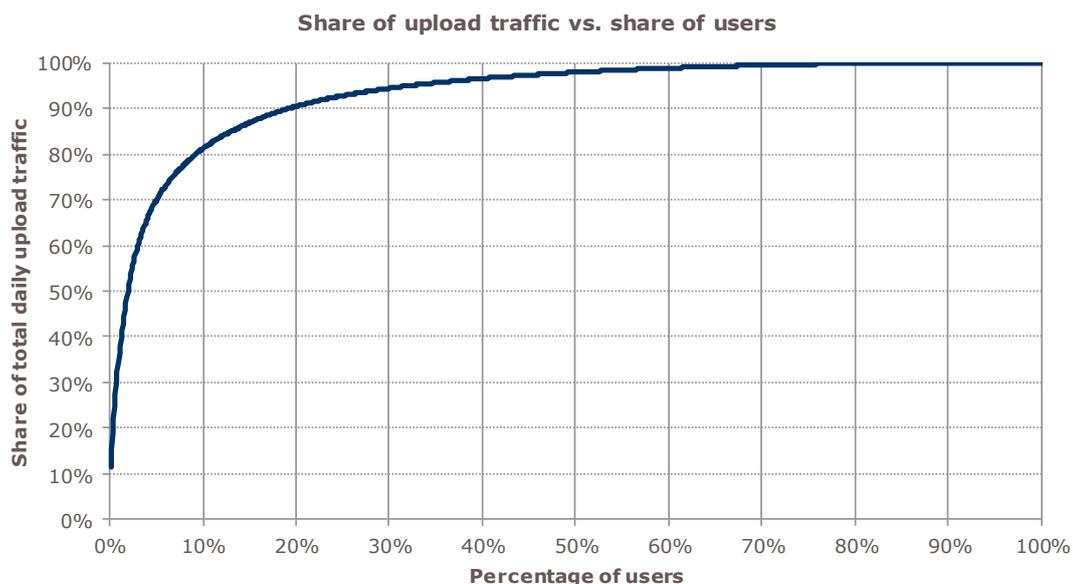


Figure 11. Distribution of ISP A subscribers to upstream traffic measured on different days in November 2013.

Using the distribution measurements, it is possible to assign traffic shares to the different user groups defined in the methodology chapter. Table 5 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. The upstream figures were derived from the above chart and Sandvine’s data, whereas the downstream figures are based on Sandvine only.

Table 5. User groups, the share of traffic they each generate and the parameters chosen for their adoption curves

	Laggards	Main	Innovators	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

Notice the different multiplication factors for upstream and downstream for each group. Due to the different factors, the ratio between upstream and downstream traffic differs between user groups. In the model, this difference is amortized over the various services. In other words, we assumed that if a user exhibits a more asymmetric or symmetric usage pattern, this is reflected in each of the services used.

Note that the power user group, although extremely small in size (2% of all users), is responsible for almost half of all upload traffic. Compared to the other groups, they also generate quite a lot of downstream traffic – although less disproportionately so than upstream traffic. By comparing the group’s traffic share and size with the average, we calculated in absolute terms how much more (or less) traffic each of the groups generates

compared to the average (this is listed in Table 5 as 'downstream/upstream traffic compared to average user'). For example, 2% of the users are power users and generate 45% of the total upstream traffic. Therefore, their total volume equals 22,5 times the average daily upstream volume.

Given the groups as shown in Table 5, 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 6.

Table 6. Estimated current demand for traffic volume by different user groups (in Mbyte/day)

	Up	Down		Up	Down
Power users	5,434	6,608	Innovators	590	2,121
Consultative web browsing	234	892	Consultative web browsing	25	286
E-mail	7	45	E-mail	1	14
Social media / Web 2.0	86	170	Social media / Web 2.0	9	55
Remote backup	561	-	Remote backup	61	-
Conversational applications	149	60	Conversational applications	16	19
Online video	395	1,953	Online video	43	627
Remote workplace	392	1,566	Remote workplace	43	503
Online music	112	135	Online music	12	43
File downloads	-	93	File downloads	-	30
Online gaming	0	1	Online gaming	0	0
Peer-to-peer file sharing	1,366	459	Peer-to-peer file sharing	148	147
Personal cloud storage	14	13	Personal cloud storage	0	0
Other services	393	765	Other services	44	250
Overhead	1,723	456	Overhead	187	146
Mainstream users	40	355	Laggards	12	37
Consultative web browsing	2	48	Consultative web browsing	1	5
E-mail	0	2	E-mail	0	0
Social media / Web 2.0	1	9	Social media / Web 2.0	0	0
Remote backup	4	-	Remote backup	0	-
Conversational applications	1	3	Conversational applications	0	0
Online video	3	90	Online video	0	0
Remote workplace	3	84	Remote workplace	1	7
Online music	0	1	Online music	0	0
File downloads	-	5	File downloads	-	0
Online gaming	0	0	Online gaming	0	0
Peer-to-peer file sharing	10	25	Peer-to-peer file sharing	-	-
Personal cloud storage	0	0	Personal cloud storage	0	0
Other services	-	22	Other services	10	22
Overhead	19	65	Overhead	1	2

5 Demand growth

This chapter discusses the growth factors in the demand for bandwidth. The structure is similar to that of chapter 4: we first distinguish the growth of different services. By adding a time dimension, the step from growth in traffic volume demand to bandwidth demand is then made. Finally, the demand growth is differentiated into various user groups.

5.1 Growth by service

At the level of services, growth is assumed to stem from two main drivers, as we explained in chapter 3. The first is growth by adoption: as more internet users start using a particular service, more traffic will be required for this service. The second driver is growth by increased intensity: as the quality of services increases, so does their traffic. The following paragraphs elaborate on the estimation parameters for growth by service.

5.1.1 Growth by adoption of services

To demonstrate and predict the adoption of services, we created a model based on historical data on adoption of services, using the current adoption of service categories in chapter 4 and the age of the first-of-its-kind service in that category. Note that adoption should be calculated as the adoption in the group of users for which the service is in any way relevant. The adoption data is therefore calculated as the adoption among the group of *potential* adopters (e.g. users for which the service is at all relevant). The percentages are estimates, primarily based on the number of people in a household thought to be using a particular service.⁷

⁷ For example, one person per household is expected to use remote workplace facilities. As the average household size is 2.3 in Western Europe (according to Eurostat, 2013), the service is relevant to 37% of all internet users.

Table 7. Age and current adoption of various internet-based service categories, used as data points for estimating the adoption curve⁸

Service	First consumer implementation	Year introduced	Current adoption	Adoption in relevant group	Modeled adoption
Consultative web browsing	WWW	1989	99%	99%	100%
E-mail	Sendmail	1983	97%	99%	100%
Social media / Web 2.0	Geocities	1994	60%	80%	96%
Remote backup	Carbonite	2004	50%	56%	15%
Conversational applications	NetMeeting	1995	66%	66%	93%
Online video	YouTube	2005	50%	53%	10%
Remote workplace	Citrix	1995	30%	86%	93%
Online music	MP3.com	1997	45%	47%	84%
File downloads	WWW	1989	83%	83%	100%
Online gaming	Ultima Online	1997	66%	88%	84%
Personal cloud storage	Dropbox	2007	11%	11%	4%

The data points are shown in Table 7. The services that fit the model are the same as found earlier, with the exception of peer-to-peer file sharing (which is treated differently, as discussed below). The 'current adoption' refers to the percentage of all internet users that has adopted a particular service. Because a service may not be relevant to all users, we therefore calculated a corrected measurement of adoption that indicates the number of users for whom the service is relevant that have adopted a service. The S-curve was matched using the latter measurement. Figure 12 shows the S-curve that best matches the data points (the last column in Table 7 shows the adoption as modeled by this curve for each service; the difference between the actual and modeled adoption is the model error).

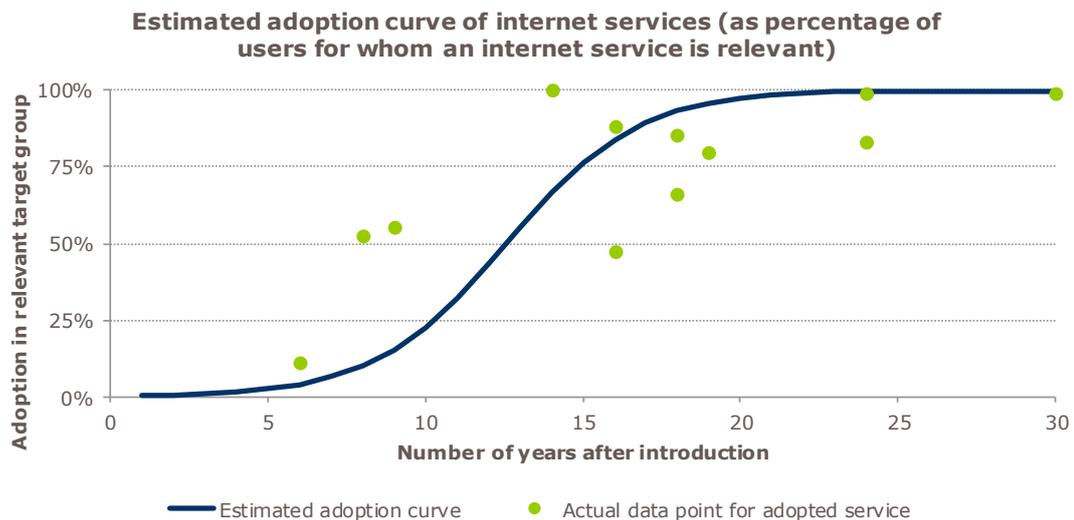


Figure 12. The adoption curve used in the model to estimate service adoption

⁸ The figures shown here were obtained from the public websites of each service and various academic and online sources estimating penetration. In some cases, the researchers estimated penetration based on input from the interviews.

Using the S-curve, it is possible to estimate the adoption rates for the coming years. Figure 13 gives an overview of the projected adoption rates per service up to the year 2020 .

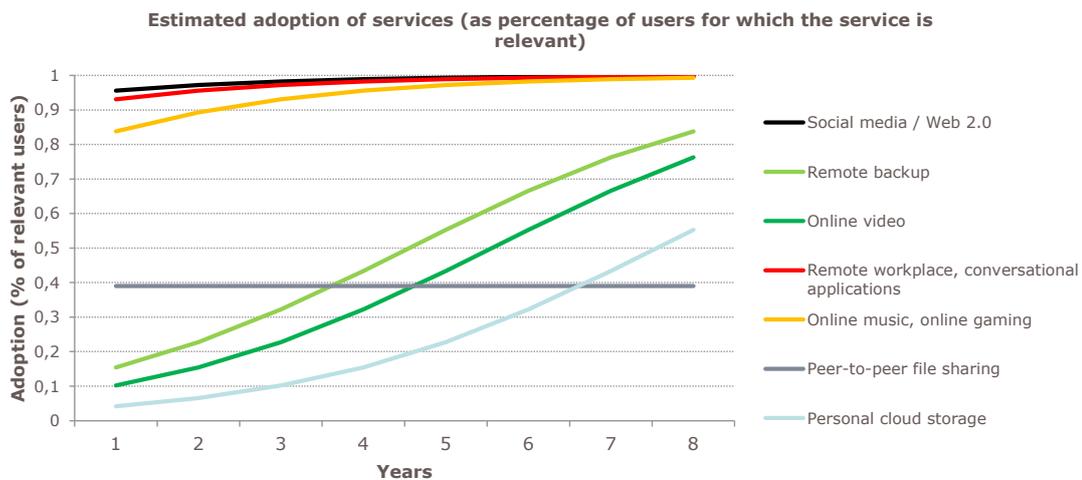


Figure 13. Estimated adoption of services used in the model, based on the S-curve

Peer-to-peer file sharing

Regarding peer-to-peer file sharing, we found that other factors influence the adoption in a non-negligible way. Peer-to-peer file sharing is not in itself a service to end users, but rather provided over the peer-to-peer network (e.g. video, audio or software downloads). The key advantage of peer-to-peer file sharing for the average end user over other services is that most of the content is illegal and free. When peer-to-peer is used for legal services, the choice whether to use peer-to-peer or centralized delivery is technical and does not involve the end user at all.

The illegality of most file transfers over a peer-to-peer network is an important factor to consider when modeling the adoption of peer-to-peer file sharing. Its adoption can be greatly influenced by the penalties for downloading or sharing illegal content or even a complete ban of such services in a country. 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 In the model, we chose to fix the adoption of peer-to-peer at the current adoption rate. The rationale behind this is that the legal alternatives are likely to be more attractive to new users than power users, who are already using peer-to-peer file sharing services. On the other hand, power users are unlikely to stop using the peer-to-peer file sharing services (although the intensity of use may drop).

5.1.2 Growth by increasing intensity

In addition to adoption of services, an increased demand for bandwidth can also be attributed to an increase in the *intensity* of services already in use. Intensity includes all aspects related to the quality of the service consumed. Part of the growth in traffic is

⁹ See for instance [9].

caused by the increased bandwidth requirements resulting from technical quality-related aspects for various types of content:

- 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 2010 was approximately 770 KB. Currently, the total average transfer size of a web site is about 1590 KB, which indicates a compound annual growth rate of 27.5% [1].
- For *software*, it seems reasonable to assume a year-over-year growth of 8% in the file size of software downloads and 10% for mobile OS updates [4]. The current size of a typical software package (e.g. a game) is assumed to be 7GB at most (this coincides with the size of a double layer DVD, which is a common medium for physical game distribution).
- For *video*, an estimate of intensity growth can be made by looking at the resolutions commonly used nowadays and the resolutions expected in a few years. An HD video stream, which is currently commonly provided through services such as Netflix and YouTube, requires around 3.8 Mbit/s of bandwidth. These video streams will probably be replaced with 4K resolution streams, which require around 15 Mbit/s of bandwidth. Assuming that 4K has replaced a fifth of the HD streams by 2023, the year-over-year growth rate in bandwidth required for video services is around 4%.¹⁰
- For *music*, there does not seem to be a demand for higher-quality audio than currently offered by most streaming services and online audio stores. The MP3 file format has been the de facto standard for music distribution for the past decade and will probably remain popular in the coming years. The MP3 file format has a maximum bit rate of 320 Kbit/s, at which the sound is virtually as good as a CD. A potential area of expansion is surround sound (e.g. more audio channels), although the additional bandwidth requirements are probably insignificant.¹¹
- For *photos*, the main driver of traffic is the resolution of images and their compression technology. Manufacturers of digital cameras traditionally competed over the number of megapixels of images, but are starting to focus more on other attributes (such as low-light performance and truthful color reproduction). Current resolutions of standard digital cameras are already more than sufficient, even for high-quality prints. We therefore do not expect much growth in traffic requirements for photographic content.
- For *online gaming*, the traffic required for sending and receiving game state information and changes therein (e.g. the traffic required to play multiplayer online games) is rather limited, and there is no reason to assume that it will grow significantly. The content, however, has grown, approximately at the same speed as video/photo material. Content includes game maps as well as game objects (e.g. a building or person in Second Life). Intensity increases as more and more games

¹⁰ The average bit rates for video streams were provided by [2].

¹¹ Surround sound can be compressed relatively well; while there are many audio channels, they are highly correlated, and modern audio compression techniques take this characteristic into account.

start downloading this content while playing.¹² As the amount of control traffic is negligible, it is safe to assume the growth rate for video for online gaming.

The intensity growth of a service is derived from the growth of what is the commonly expected quality for photo, video, music, software or web page size according to the above aspects. Services not accounted for ('other services') are assumed to grow with the average growth rate in the model (24% year-over-year).

Peer-to-peer file sharing

In paragraph 5.1.1 we elaborated on how peer-to-peer file sharing is expected to follow a different adoption pattern. Regarding the intensity growth, we also expect peer-to-peer file sharing to follow a different route than the rest of the service categories. As can be seen in Figure 8, currently almost 70% of the upload traffic is comprised of traffic from peer-to-peer and newsgroups. ISPs indicate that the typical usage pattern is that a small group of extreme users take up all available capacity, regardless of the demand. This usage thus does not seem to reflect demand but rather supply-driven bandwidth consumption.

As indicated, the current model assumes the bandwidth consumed is driven by demand rather than supply. However, given the fact that peer-to-peer file sharing has such a large impact on traffic, we have chosen to incorporate this specific behavior in the model. We assume that the growth of peer-to-peer traffic follows the growth of provisioned speeds. Calibrated on the increase of the provisioned speed of ISP A over the year 2013, we have modelled the year-over-year growth of peer-to-peer file sharing to be 40%.

An increased take-up of other streaming protocols, such as peer-to-peer for video, could be potentially disruptive to bandwidth usage. An example of a currently popular application is SopCast, which delivers video stream using peer-to-peer technology, and is typically used for watching live sports events. The interviewed experts however do not expect peer-to-peer video to take up, especially given the fact that legal over-the-top alternatives are coming about rapidly. The experts additionally argued that the take-up of such services is limited by the fact that quality of service is harder to guarantee for such services.

5.2 Growth in capacity and speed demand

One of the ways demand can grow is when the access connection is used to consume more at the same time. Although the bandwidth required by a single service does not change, the capacity needed to accommodate multiple services at one time is much higher. One obvious cause could be that the same connection is being shared by more users in a home than before. Nevertheless, in the time period covered by this study, a significant growth in the average household size is not expected. Eurostat estimates this size to be 2.3 persons per household [14]. In addition, the average age at which children go online is not expected to decrease significantly.

Another cause of growth from simultaneous usage could be that users start using more than one service at the same time. This trend was confirmed in our qualitative study. A common scenario is that people start using so-called 'second screen' applications on their tablet, while watching a video stream on their television. In the model, we assume that such 'secondary' applications are not as bandwidth-intensive as the primary application. The underlying assumption is that people simply cannot consume more than a certain amount of content in a given time period.

¹² Also compare to [16]

Simultaneous use can also be 'automatic' in the sense that services use bandwidth in the background together with other services. In the future, automatic software updates will become more and more common, requiring less or no user intervention. On the other hand, several up-and-coming services such as remote back-up and cloud storage are not urgent and can spread their traffic over the course of a whole day. Whether this actually happens may still depend on the user, who is for instance likely to turn off his computer and other devices during the night (some services will continue to upload data even then – the automatic backup from an iPhone by default runs at night, when the phone is not in use and connected to its charger).

5.3 Growth from future revolutionary services

Up until this point, we have discussed the development of currently existing services. However, the historical development of internet services has shown us 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.

Even though the moment such services are introduced is impossible to predict, it is feasible to model the effect of such services on bandwidth demand by formulating several scenarios. In each scenario, we assume a particular frequency with which a 'revolutionary' service is introduced and a particular combination of average upstream and downstream traffic. The adoption growth for a future service can be modeled using the fitted S-curve (assuming the future service is relevant for all internet users), which seems to model historical 'revolutionary' services pretty well.

For the purposes of this study, we calculated the expected probability of the development of future services and their expected impact on demand growth, as described in paragraph 3.7.1. Table 8 is an excerpt of a list we used to model revolutions. Each row represents a particular kind of revolution: a 'once every year' revolution occurs frequently, but has a negligible impact on growth; the 'once every fifty years' kind of revolution is rare, but has enormous impact. For each type, we calculated the expected number of times such a revolution will occur in the seven-year period. We subsequently modeled 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.

Table 8. Excerpt from the table modeling the expected growth in bandwidth in a seven-year period.

Revolutionary service type	Next seven years		Estimated impact on traffic demand		Expected growth in seven year period	
	Chance of at least one introduction	Expected number of services	Min	Max	Min	Max
Once every 1 year	100%	7,00	0,00%	0,00%	0,0%	0,0%
Once every 2 years	99%	3,50	0,00%	0,00%	0,0%	0,0%
...
Once every 49 years	13%	0,14	23,06%	55,34%	3,3%	7,9%
Once every 50 years	13%	0,14	25,00%	60,00%	3,5%	8,4%

Figure 14 gives a graphical overview of the distribution of the probability of revolutionary events happening and the corresponding impact of such an event on bandwidth demand.

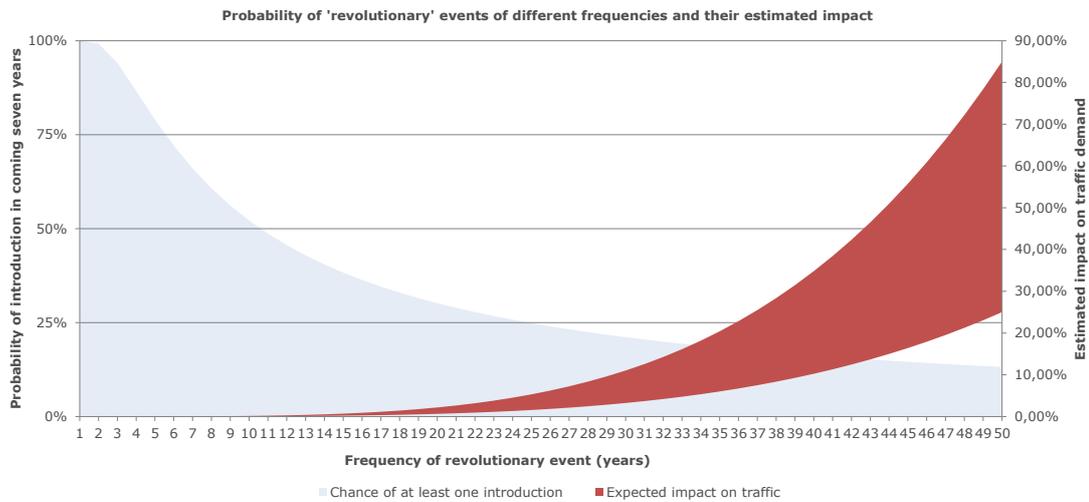


Figure 14. Probability of 'revolutionary' events of different frequencies and their estimated impact

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.

5.4 Differences between user groups

Figure 12 shows the adoption curve estimated for services. 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 15 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.

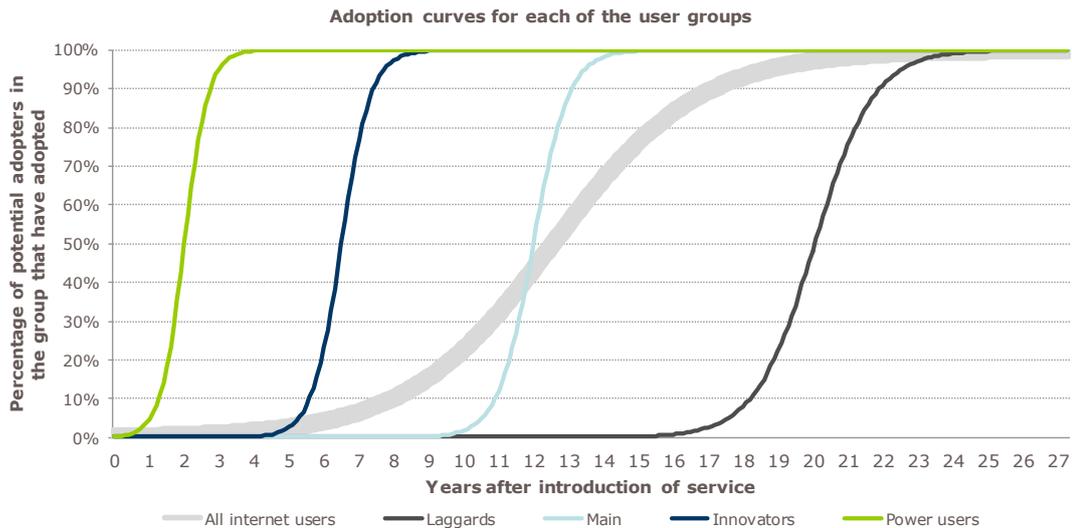


Figure 15. The adoption curves defined for each of the four user groups that add up to the total adoption curve used earlier

Using the different adoption curves for each user group, it is possible to calculate different growth factors per group for each service. Figure 16 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).

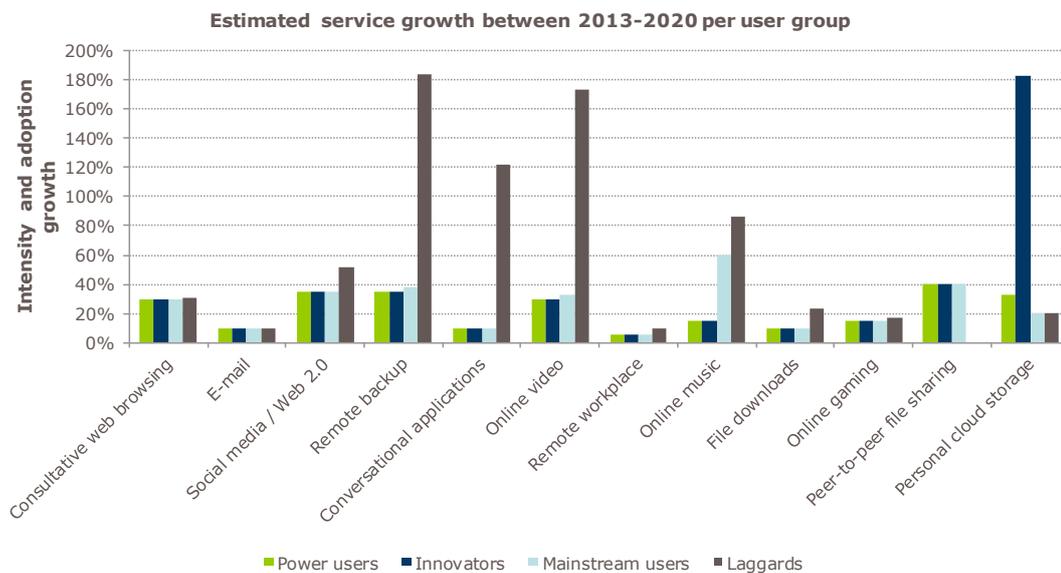


Figure 16. Growth rate of services for each user group

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.

6 Future demand

In this chapter we discuss the future demand for bandwidth. As in section 4, we started with the aggregate demand and added further detail to distinguish the various services. Including a time dimension enabled us to make the step from traffic volume demand to bandwidth demand.

6.1 Aggregate demand

The model showed an estimated CAGR of 40.3% for downstream and 43.9% for upstream traffic. Figure 17 gives an overview of the resulting projected traffic volumes for the coming seven years relative to the traffic volume in 2013. Note that this increase includes growth from future revolutionary services. The growth of existing services alone is responsible for a CAGR of 36.6% (upstream) and 31.6% (downstream) respectively.

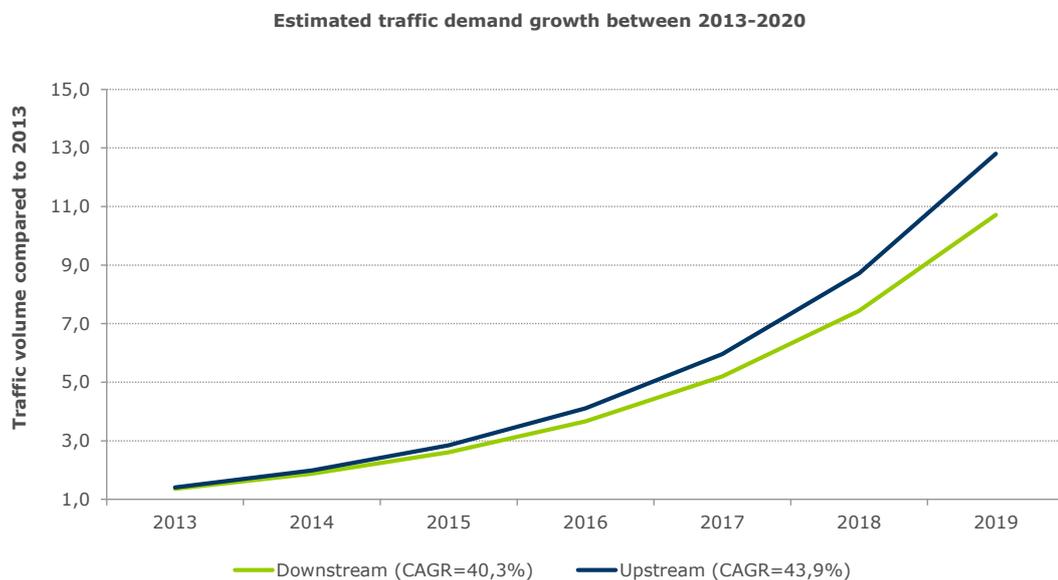


Figure 17. Estimated growth in traffic volume demand relative to 2013 according to the model

6.2 Demand by service

Figure 18 and Figure 19 show the estimated volumes for upload and download respectively generated by existing services, grouped by the different service types.

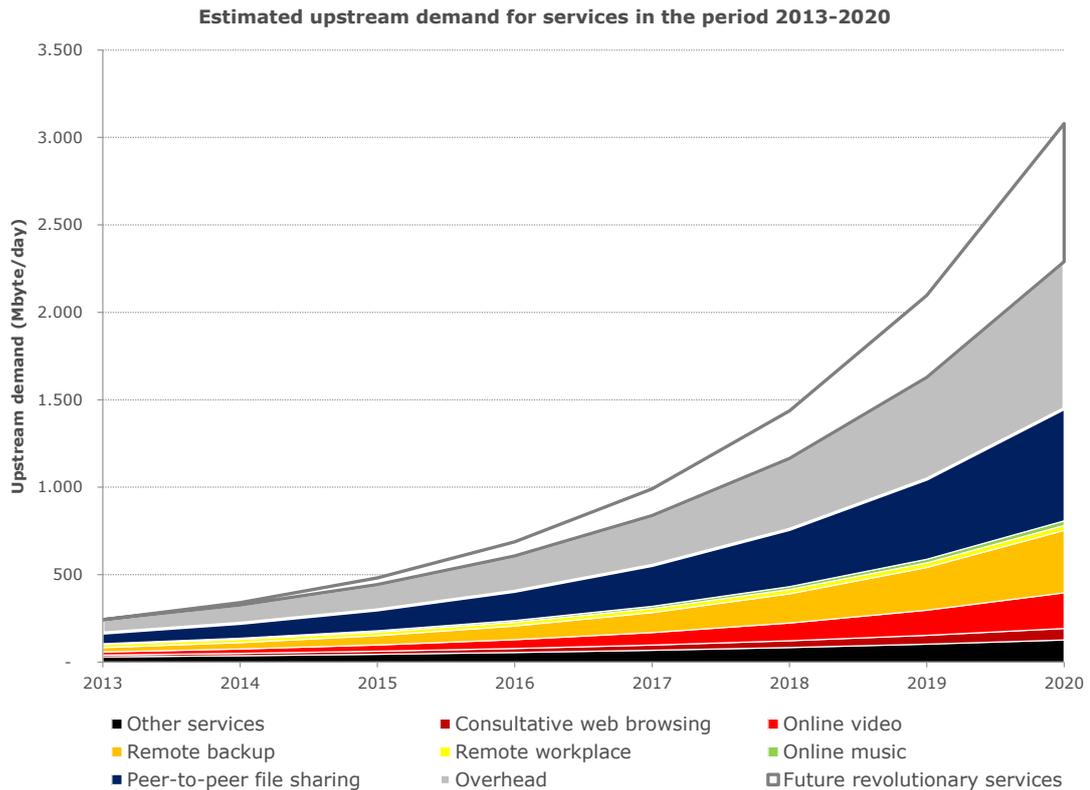


Figure 18. Projected average daily upstream traffic volume per residential subscription for the years 2013-2020

Regarding the upstream demand, the model shows that in addition to future revolutionary services, overhead traffic and peer-to-peer file sharing will continue to comprise the majority of the upload traffic. Since overhead traffic partly consists of acknowledgements of the downstream traffic, this demand is mainly driven by a high download demand. We modeled peer-to-peer traffic to only grow in intensity and not adoption, and is moreover modeled to be supply-driven rather than demand-driven. It is therefore subject to the growth of the upstream capacity provided by operators. Other demanding services in 2020 are remote backup services, online video and future revolutionary services. The intensity growth for consultative web browsing may be caused by web sites moving from the 'consultative' to the more 'web 2.0' type of service (which is not modeled). Also, the fact that video services are driving a significant share of upstream traffic is food for further analysis. User-generated video was not included in the definition of this service, so the upstream video traffic probably mostly result from current peer-to-peer distribution models being used for video. The estimated demand for daily upstream traffic in 2020 will average at just over 3,000 Mbyte per day.

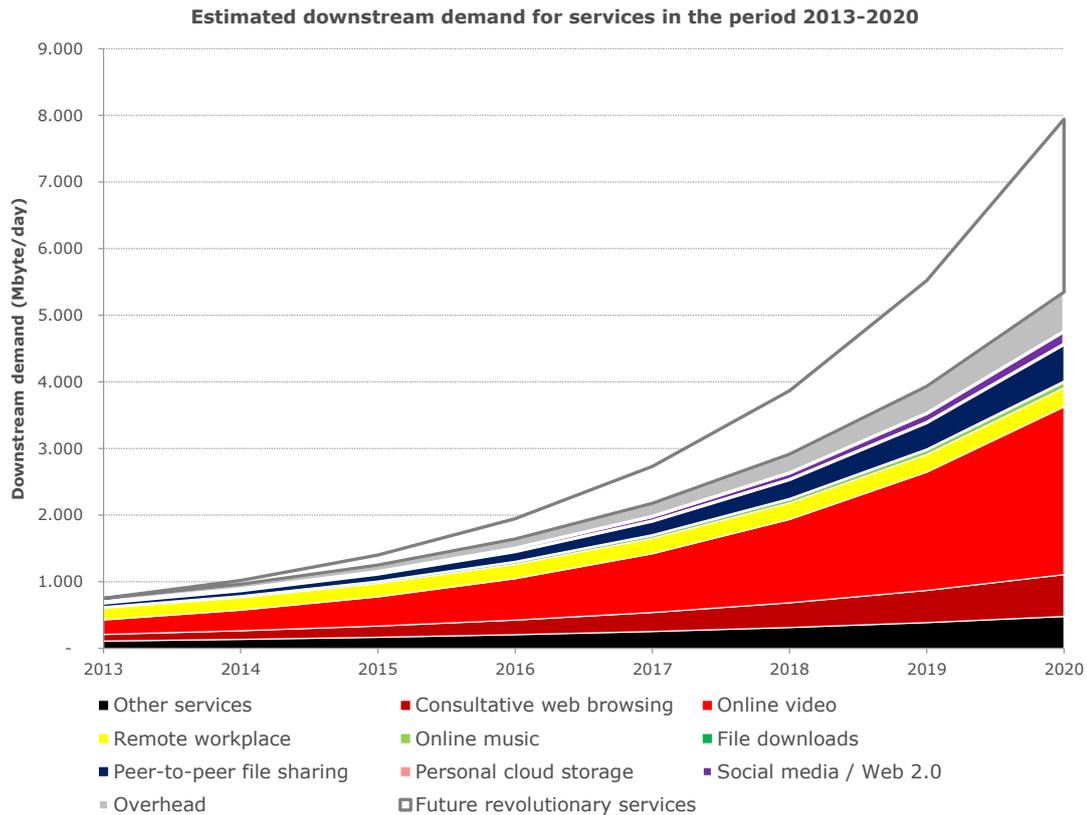


Figure 19. Projected average daily downstream traffic volume per residential subscription for the years 2013-2020

Unsurprisingly, online video is a major driver of downstream traffic growth. Besides online video, traffic for consultative web browsing, overhead and peer-to-peer file sharing are also expected to require considerable traffic. Again, the future revolutionary services also comprise a large part of the traffic in 2020. It is interesting to see web browsing pop up as one of the bigger drivers of growth in this model. The increase in web browsing is mainly caused by intensity growth (e.g. increasing web page size). The growth of online video consumption is due to intensity growth (i.e. the move towards HD and higher resolutions) as well as growth from adoption by lagging users. Other drivers of downstream traffic growth are overhead and peer-to-peer file sharing. Growth in traffic for peer-to-peer file sharing is solely driven by intensity growth. The total downstream demand for 2020 is estimated at almost 8,000 Mbyte per day.

6.3 Demand for capacity and speed

As discussed earlier, the amount of *bandwidth* required depends on the volume of the traffic to be transferred on a given day and 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 bandwidth to end-users than the minimum bandwidth 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 a bandwidth 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 bandwidth required to satisfy the demand is 4.2 Mbit/s. In that case, the ISP will 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 estimated the current urgency of demand by comparing the current daily traffic volume with the currently provisioned connection speeds by ISP C. We find that on average, 80% of a user's daily downstream traffic is transferred in less than 6 minutes. Regarding the upstream traffic, users on average transfer 80% of a day's traffic in 17 minutes.

Assuming that the urgency of the demand in 2020 will not be different from the urgency in 2013, it is possible to calculate the sufficient provisioned speed in 2020 given the traffic volume predicted. If the user in the example above has a demand of 1 gigabyte per day in 2020, the minimum sufficient provisioned speed is 21.8 Mbit/s (the connection needs to be sufficiently fast in order to transfer 80% of 1 gigabyte in five minutes).

As an example, consider a user in the mainstream group. In 2020, this user is expected to download (on average) 2.6 gigabyte of traffic and upload 547 Mbyte of traffic per day. Assuming that this user still wants to transfer 80% of his downstream traffic in six minutes or less, the sufficient provisioned downstream speed would be 55 Mbit/s. Following the same line of reasoning, the sufficient provisioned upstream speed for this user in 2020 would be 3.5 Mbit/s.

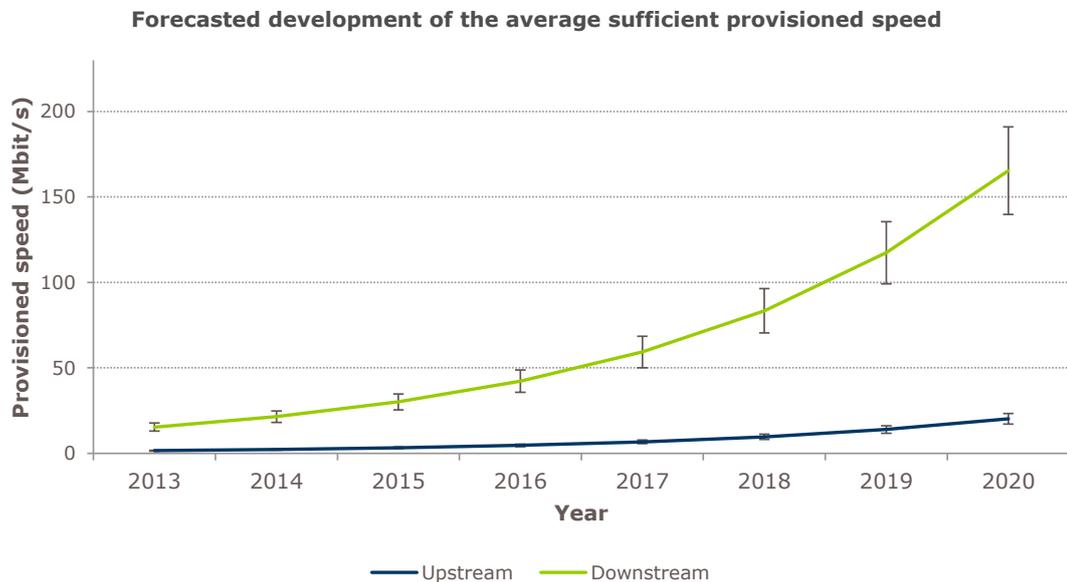


Figure 20. Estimated development of the average sufficient provisioned speed of subscriptions

Figure 20 shows the forecasted development of the average sufficient provisioned subscription speeds using the method described. You can see from the graph that an average subscription will have a sufficient provisioned downstream speed of about 165.4 Mbit/s in 2020 (compared to 15.3 Mbit/s in 2013) and an average sufficient provisioned upstream speed of 20.1 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 20 show the speeds required if urgency changes by 20% (i.e. traffic needs to be transferred in 20% more or less time than currently).

6.4 Differences between user groups

6.4.1 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 21 and Figure 22 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.

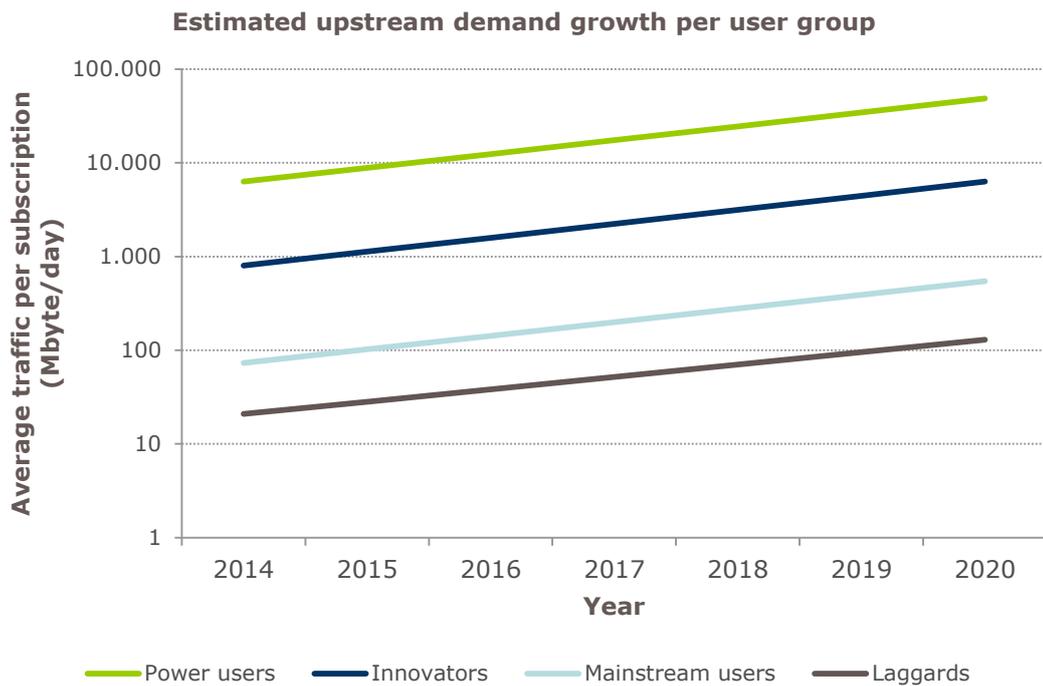


Figure 21. Growth in upstream demand per user group as estimated by the model

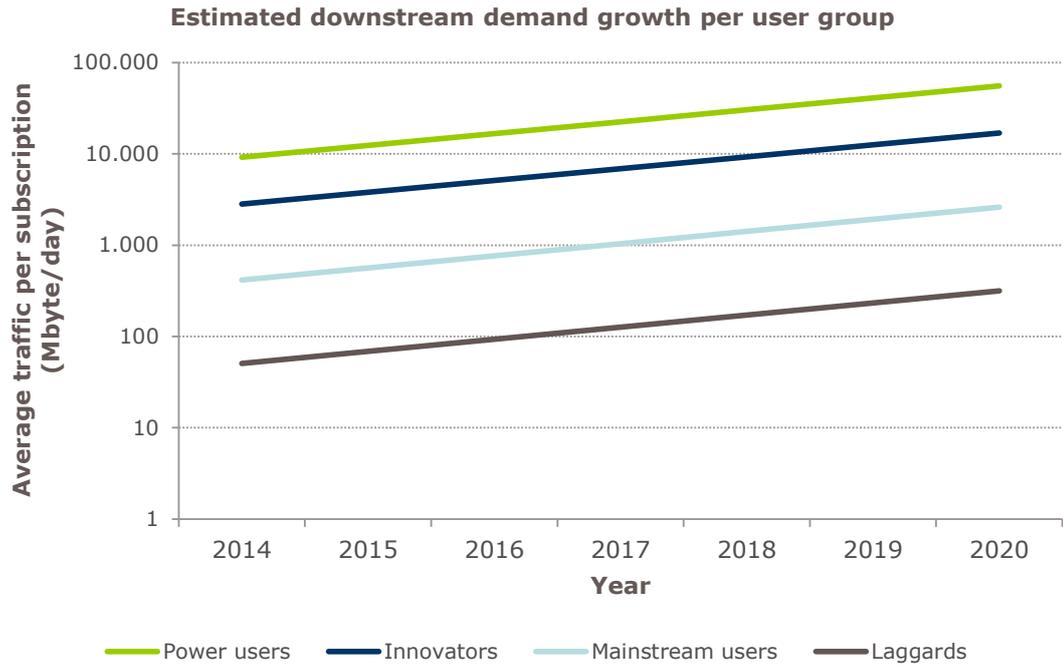


Figure 22. Growth in downstream demand per user group as estimated by the model

6.4.2 Demand by service per user group

Table 9 shows the estimated traffic volume demand per subscription in 2020 for the various user groups, broken down into different service categories.

Table 9 Estimated bandwidth demand in 2020 (in Mbyte/day transferred) for an average subscription in the user group, by service

	Up	Down		Up	Down
Power users	48,704	55,396	Innovators	6,312	16,924
Consultative web browsing	1,466	5,597	Consultative web browsing	159	1,797
E-mail	14	87	E-mail	2	28
Social media / Web 2.0	707	1,390	Social media / Web 2.0	77	446
Remote backup	4,585	0	Remote backup	498	0
Conversational applications	290	116	Conversational applications	32	37
Online video	2,479	12,254	Online video	269	3,933
Remote workplace	551	2,204	Remote workplace	60	707
Online music	299	358	Online music	32	115
File downloads	0	181	File downloads	0	58
Online gaming	1	3	Online gaming	0	1
Peer-to-peer file sharing	14,403	4,835	Peer-to-peer file sharing	1,565	1,552
Personal cloud storage	100	96	Personal cloud storage	11	31
Other services	1,773	3,448	Other services	199	1,126
Overhead	6,965	6,613	Overhead	1,795	1,236
Future revolutionary services	15,071	18,212	Future revolutionary services	1,614	5,856
Mainstream users	547	2,619	Laggards	129	317
Consultative web browsing	11	301	Consultative web browsing	3	31
E-mail	0	5	E-mail	0	0
Social media / Web 2.0	5	75	Social media / Web 2.0	2	8
Remote backup	34	0	Remote backup	5	0
Conversational applications	2	6	Conversational applications	1	1
Online video	18	658	Online video	3	31
Remote workplace	4	118	Remote workplace	1	12
Online music	2	19	Online music	0	0
File downloads	0	10	File downloads	0	1
Online gaming	0	0	Online gaming	0	0
Peer-to-peer file sharing	107	260	Peer-to-peer file sharing		
Personal cloud storage	0	0	Personal cloud storage	0	0
Other services	0	101	Other services	46	100
Overhead	261	152	Overhead	34	124
Future revolutionary services	102	915	Future revolutionary services	35	108

6.4.3 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. Estimated development of average sufficient provisioned speeds (in Mbit/s) for different user groups (2013-2020)

		2013	2014	2015	2016	2017	2018	2019	2020
Power users*	Up	29.2	40.8	57.2	80.2	112.7	158.6	223.4	314.9
	Down	142.4	191.4	257.5	346.9	467.9	631.6	853.5	1154.5
* Note power users:		The estimations for the sufficient provisioned speeds for <u>power users</u> are based on a different method in which traffic for peer-to-peer is modeled to be supply-driven rather than demand-driven. This means that the power users will always maximally utilize the provisioned bandwidth.							
Innovators	Up	3.7	5.2	7.3	10.3	14.4	20.4	28.8	40.8
	Down	44.0	59.0	79.3	106.7	143.7	193.6	261.2	352.7
Mainstream users	Up	0.3	0.5	0.7	0.9	1.3	1.8	2.5	3.5
	Down	6.4	8.7	11.7	15.9	21.6	29.4	40.1	54.6
Laggards	Up	0.1	0.1	0.2	0.2	0.3	0.5	0.6	0.8
	Down	0.8	1.1	1.4	1.9	2.6	3.6	4.9	6.6
All users	Up	1.6	2.2	3.2	4.6	6.7	9.6	13.9	20.1
	Down	15.3	21.4	30.0	42.2	59.3	83.4	117.4	165.4

7 Discussion & conclusions

In this chapter, we return to the research questions posed in chapter 1 and show how they can be answered using the bandwidth demand estimation model. We subsequently discuss some limitations of the model.

7.1 Answering the research questions

The central question in this study is how the demand for upstream traffic will develop in the coming years until the year 2020. For that purpose, we constructed a bandwidth demand estimation model. To do this, we took the following steps: first, we estimated the current demand for services. Second, we estimated the growth of (a) the adoption and (b) the intensity of these services. Moreover, we added 'future revolutionary services', services that are by definition not in use yet but are expected to come about in the coming years. These estimates led to a demand for megabytes per day. At the last step, we translated this demand for volume to demand for bandwidth speed, based on the urgency of the traffic. Upload traffic in essence appeared to be less urgent than download speed, as was shown by the network measurements conducted.

The services that contribute most to current upstream traffic are peer-to-peer file sharing (61 Mbyte), remote back-up (25 Mbyte) and online video (18 Mbyte). These services are in 2020 also expected to be the largest driver for upstream traffic. The main driver for downstream traffic is online video, comprising amongst others traffic from YouTube and Netflix. Remote workplace, consultative web browsing and again peer-to-peer file sharing are important drivers for downstream demand.

It is important to recognize that residential broadband speed demand has a very diverse character. Households vary considerably in their intensity of use, the type of applications they use their connection for, and the amount of traffic these applications generate. In this study, we have addressed this diversity by differentiating between different categories of users, and our overall outcomes are averages over these categories.

Six subquestions were to be answered by this research. We will describe the answers separately.

1. To what extent do currently available applications contribute to upstream traffic by 2020?

47% of the upstream bandwidth traffic demand in 2020 is constituted by the currently available applications. Overhead traffic and peer-to-peer file sharing comprise the majority of the upload traffic, constituting respectively 27% and 21% of the upstream traffic demand in 2020. Other currently available applications that drive demand in 2020 are remote backup (12%) and online video (7%).

2. To what extent has the need for upstream traffic of currently available applications changed in recent years?

The model estimates that upstream traffic in 2013 is 242 Mbyte per day. By 2020, upstream is expected to take up 3,093 Mbyte per day. This constitutes a CAGR of 44%. The CAGR of the upstream traffic of currently available applications only is 37%. The largest CAGR is expected to come from remote backup (46%) and online video (42%).

In the past few years, Cisco VNI [7] as well as TNO and Dialogic [8] conducted studies to estimate the growth of bandwidth consumption at the aggregate level. Cisco reports that it expected a compound annual growth rate (CAGR) of global IP traffic of 32% for the period 2010-2015, including fixed as well as mobile access. In 2010, TNO and Dialogic estimated a growth rate of 30 to 40% in the Netherlands for fixed-only connections.

3. Which business applications require high upstream traffic?

We find that the current bandwidth demand for business applications suitable for use in a residential context (working from home) is low. Applications that require large amounts of upstream bandwidth are typically applications that either transmit large files or multimedia content, or use peer-to-peer as distribution mechanism. We were unable to identify significant bandwidth usage following either pattern for business applications in general. It should be noted that our model estimates the average demand for a large group of users, and cannot take into account the requirements for highly specific business applications.

We conclude that bandwidth demand for business applications mainly stems from simultaneous use of multiple applications, or simultaneous use by multiple users. In addition, the use of consumer applications for business purposes contributes (albeit invisibly) to the upstream traffic demand for business use.

4. To what extent will consumers use more business applications by 2020?

Services that are typically also used for business applications are conversational applications, remote workplace, and personal cloud storage, as elaborated on in the previous question. The CAGR of the demand for these services is respectively 17%, 7% and 25%. By 2020, these services will have a relatively limited impact on the traffic demand for residential connections: conversational applications take up 0.6% of the upstream and 0.2% of the downstream traffic; remote workplace service consume 0.9% of the traffic (up) and 3.6% (down). Personal cloud storage is also marginal in 2020 with 0.2% of the upstream traffic and 0.2% of the downstream traffic.

5. To what extent will consumers use other applications with a high demand for upstream traffic by 2020?

In this research, we modeled demand for future revolutionary services. Future revolutionary services are services that do not exist at this moment, but are expected to come about in the coming years. We modeled their traffic by means of a probability distribution of the impact of revolutions and their expected outcome frequency. The future revolutionary services are expected to constitute 26% of the upstream traffic by 2020. One typical example of a driver for this traffic is a surge in the number of connected devices and accompanying services in a household.

6. Which upstream and downstream speeds will be sufficient for future demand?

The average sufficient provisioned downstream bandwidth speed in 2020 is 165 Mbit/s, whereas the average sufficient provisioned upstream speed is 20 Mbit/s. Regarding the downstream demand, 59% of the traffic is comprised by currently available applications. 34% of the traffic demand is constituted by future revolutionary services, whereas 7% is comprised by overhead traffic. Regarding the upstream demand, 47% of the traffic is constituted by currently available applications. 26% of the traffic demand comes from future revolutionary services, whereas 27% is comprised by overhead traffic.

Bandwidth usage is modelled to be demand-driven. Bandwidth demand is estimated to be different between different types of users. Power users are expected to be the most demanding, followed respectively by innovators, mainstream users and laggards. However,

in the case of power users, we assume the demand for 'peer-to-peer file sharing' to be supply-driven. The intensity growth of demand for peer-to-peer services was calibrated on the increase of the provisioned speeds of ISP A in 2013.

7.2 Limitations

The model developed in this study estimates future demand for bandwidth 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 modeling 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. First of all, some traffic may be 'involuntary'. One example is the traffic generated by viruses and spyware, which recent research has shown to account for up to 61% of online traffic [15]. This could also be a reason that our estimates of current traffic differ from other studies, the European Commission's report 'Quality of Broadband Services in the EU' being a case in point. Second, 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 bandwidths of their subscriptions, they would see significant increases in traffic volume, which nevertheless do not directly relate to increased *demand*.

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:

- (Further) criminalization 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 risk of being caught, 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.

7.3 Conclusion

In this study, we have developed a method for estimating households' future bandwidth demand. The future bandwidth demand for existing services can be estimated by combining data on current usage of services with projections of the adoption and growth rates of service usage intensity. Future revolutionary services are expected to play a major role in the growth of bandwidth demand. Such services are however difficult to foresee. We have developed a 'next best' estimation model for these services, where we modeled a probability distribution of the impact of revolutions and their expected occurrence frequency.

The method developed in this study was applied to data on residential subscriptions in Western Europe. We predicted the compound annual growth rate (CAGR) of upstream and downstream traffic demand to be 44% and 40% respectively. While demand in 2013 is on average 15.3 Mbit/s downstream and 1.6 Mbit/s upstream, in 2020 demand is expected to increase to 165.4 Mbit/s downstream and 20.1 Mbit/s upstream. Large differences can be found between the types of services and the user groups. Power users, constituting 2% of the total users, will require 1,155 Mbit/s downstream and 315 Mbit/s upstream by 2020, whereas the laggards will only need 6.6 Mbit/s downstream and 0.8 Mbit/s upstream by that time.

Appendix A. Consulted experts

Name	Organization	Function
Arbanowski, Dr. ing. Stefan	Fraunhofer FOKUS	Head of Future Applications and Media
Betteridge, Ian	iPoque	International Sales Specialist
Belson, David	Akamai	Editor of 'State of the Internet' report
Coppieters, Stijn	Genexis	Marketing Manager
Osstyn, Dirk	Alcatel-Lucent Belgium	Global Network Engineering / Fixed Access Program Assurance
Pesovic, Ana	Alcatel-Lucent Belgium	Marketing Manager
	Ftth Council Europe	Chair of World Applications Committee
Poort, Joost	Institute of Information Law (IViR)	Senior researcher
Schuurman, Krijn	InterimIC	Project manager
Vries, Wilbert de	Tweakers.net	Editor-in-Chief
Ziel, Stef van der	JetStream	Owner

References

- [1] The HTTP Archive (2013). *Trends*. Retrieved from <http://httparchive.org/trends.php?s=All&minlabel=Nov+15+2010&maxlabel=Nov+15+2013>
- [2] Sandvine (2013) The Global Internet Phenomena Report. 2013 H1
- [3] Sandvine (2013) The Global Internet Phenomena Report. 2013 H2
- [4] NationMaster (2014). International Internet Bandwidth in Mbps per capita. Retrieved from <http://www.nationmaster.com/country-info/stats/Media/Internet/International-Internet-bandwidth/Mbps/Per-capita>
- [5] Kenny & Broughton (2013) Domestic demand for bandwidth: an approach to forecasting requirements for the period 2013-2023. commissioned by the Broadband Stakeholders Group (BSG)
- [6] Microsoft (2008) Remote Desktop Protocol Performance white paper. Retrieved from http://download.microsoft.com/download/4/d/9/4d9ae285-3431-4335-a86e-969e7a146d1b/RDP_Performance_WhitePaper.docx
- [7] Cisco (2011) Visual networking index (VNI) Global IP Traffic forecast 2010-2015. Presentation by Mark Nowell. Retrieved from http://www.ieee802.org/3/ad_hoc/bwa/public/sep11/nowell_01_0911.pdf
- [8] Dialogic and TNO (2010). Supply and demand for next generation infrastructures
- [9] Poort et al. (2013) Baywatch: two Approaches to Measure the Effects of Blocking Access to The Pirate Bay
- [10] Rogers, E.M. (2005) Diffusion of Innovations.
- [11] Cisco (2013). Cisco Visual Networking Index: Forecast and Methodology. 2012-2017 – white paper
- [12] Nielsen (2005). Nielsen’s Law of Internet Bandwidth. Retrieved from <http://www.nngroup.com/articles/law-of-bandwidth/>
- [13] Crozier, R. (2013). Networks hammered by Apple iOS7 release. Retrieved from <http://www.itnews.com.au/News/357682.networks-hammered-by-apple-ios7-release.aspx>
- [14] Eurostat (2013). Average household size. Retrieved from <appsso.eurostat.ec.europa.eu/nui/submitViewTableAction.do>
- [15] Incapsula (2013). Report: bot traffic is up to 61.5% of all website traffic. Retrieved from <http://www.incapsula.com/the-incapsula-blog/item/820-bot-traffic-report-2013>
- [16] Pesovic, A. & Sharpe, R. (2012). Is symmetrical bandwidth a Myth or a Must? Retrieved from <http://www2.alcatel-lucent.com/techzine/is-symmetrical-bandwidth-a-myth-or-a-must>.
- [17] ECC (2013). Future Harmonised Use of 1452-1492 MHz in CEPT. Copenhagen: ECC.



Contact:

Dialogic
Hooghiemstraplein 33-36
3514 AX Utrecht
The Netherlands
Tel. +31 (0)30 215 05 80
Fax +31 (0)30 215 05 95
www.dialogic.nl