

Quality of Broadband Services in the EU

March 2012

FINAL REPORT

A study prepared for the European Commission
DG Communications Networks, Content & Technology

*Digital
Agenda for
Europe*

This study was carried out for the European Commission by SamKnows Limited



SamKnows

Internal identification

Contract number: 30-CE-0392545/00-77
SMART 2010/0036

DISCLAIMER

By the European Commission, Directorate-General of Communications Networks, Content & Technology.

The information and views set out in this publication are those of the author(s) and do not necessarily reflect the official opinion of the Commission. The Commission does not guarantee the accuracy of the data included in this study. Neither the Commission nor any person acting on the Commission's behalf may be held responsible for the use which may be made of the information contained therein.

Copyright © 2013 – European Union. All rights reserved. Certain parts are licensed under conditions to the EU.

Reproduction is authorised provided the source is acknowledged.

[THIS PAGE LEFT INTENTIONALLY BLANK]

Contents

A	EXECUTIVE SUMMARY	7
A.1	Background	7
A.1.1	Purpose of the study	7
A.1.2	About SamKnows	9
A.1.3	Overview of Methodology	9
A.2	Key Findings	12
A.2.1	Summary	12
A.2.2	Performance by Technology	12
A.2.3	Performance by State	13
A.3	Next Steps	14
B	METHODOLOGY AND DEFINITIONS	16
B.1	Methodology	16
B.1.1	An Open and Transparent Test Methodology	16
B.1.2	Measurement Methodology	17
B.1.3	Sample Plan Methodology	39
B.1.4	Final Panel Composition	43
B.1.5	Data Analysis Methodology	49
B.2	Definitions	51
B.2.1	Technology Splits	51
B.2.2	Data Dictionary	53
B.3	Further information regarding sampling methodology	56
B.3.1	Identifying the Test Variables	56
B.3.2	Subdivisions of the Population	56
B.3.3	Sample Parameters	57
B.3.4	Participant Recruitment	57
B.3.5	Statistical Analysis Considerations	58
B.3.6	Data Validation	59
B.3.7	Impact of Unaccounted Variables (IPTV)	60
C	EU LEVEL ANALYSIS	61
C.1	Key Performance Indicators	61
C.1.1	Download Speed	61
C.1.2	Actual Download Speed split by hour of day and technology	63
C.1.3	Upload Speed	64
C.1.4	Latency	66

C.1.5	Packet Loss	68
C.1.6	DNS Resolution	70
C.1.7	Web Browsing	74
C.1.8	VOIP	77
C.2	Comparison with the United States	81
C.2.1	Download Speed	81
C.2.2	Upload Speed	81
C.2.3	Latency	82
C.2.4	Packet Loss	82
D	COMPARISON BETWEEN COUNTRIES	83
D.1	Key Performance Indicators	83
D.2	Download and Upload Speeds	83
D.2.1	Download	83
D.2.2	Upload	87
D.2.3	Latency	92
D.2.4	Packet Loss	95
D.2.5	DNS Resolution Time and Failure Rate	97
D.2.6	Web Browsing Speeds	101
D.2.7	VoIP Jitter	103

About This Document

This document is the first in a series of reports, which have been commissioned by the European Commission and will be completed by SamKnows over a three year period. For the purposes of this study 9,104 households across the European Union were given a specially configured hardware device (SamKnows Whitebox), which runs a series of purpose-built tests to measure every aspect of Internet performance.

Together the document provides a comprehensive explanation of the project, the purpose, the test methodology, analysis of performance against key indicators across the EU and individual country-by-country analysis.

This first study has been carried out without the assistance of European ISPs, but it is hoped that in future studies, ISPs will participate on a level that is equal to the contributions made by ISPs in the SamKnows/FCC/USA Measuring Broadband America Project.

The analysis in this report is carried out on data collected in the month of March 2012. The project is ongoing and SamKnows continues to look for volunteers to participate in the study by signing up at <http://www.samknows.eu/>

Any comments on the analysis in this document should be directed to SamKnows at team@samknows.com

Executive Summary

Background

Purpose of the study

In March 2010 the European Commission adopted “Europe 2020”, a strategy for European economic and social development to promote smart, sustainable and inclusive growth to stimulate a high-employment economy to deliver social and territorial cohesion throughout the Member States. A key part of this initiative is a target to achieve universal broadband access by 2013 and give citizens access to much faster internet speeds across Europe by 2020. Higher broadband speeds have been defined as 30 Mbps or above, with a further goal of 50% or more European households subscribing to broadband connections above 100 Mbps.

This study falls under the “Digital Agenda for Europe” which was adopted on 19 May 2010. The focus of this Agenda is a framework for stimulating growth and innovation notably through maximizing the potential of Information and communication technologies (ICTs). This initiative builds on previous activity by the European Commission, which has been monitoring take-up of broadband access in the EU since July 2002 through the Communications Committee. This research has shown that whilst progress has been made in extending broadband coverage, with 95.7% of Europeans now able to access high speed broadband, the figure drops to 78% in rural areas, and in some countries high speed broadband cover just 60% or less of the rural population.

The European Commission therefore commissioned a study on broadband performance to obtain reliable and accurate statistics of broadband performance across the different EU Member States and other countries. This data will be necessary for the benchmarking of the European Digital Agenda, the European Initiative for the development of the information Society. Given the critical nature of this data, it is imperative that the methodology used be open and transparent and that data be made available for comprehensive review.

After an extensive international tender, SamKnows was selected to carry out the study and this report is the first in a series of documents, which will be released over the course of the three year study.

To undertake this study, SamKnows used its now globally accepted methodology which is also used by governments and regulators in Europe, North America, South America and Asia to measure both fixed and mobile internet performance. This report focuses solely on European fixed broadband performance and compares the advertised speed against the effective broadband speed, as experienced by the consumer.

Previous studies have shown that the effective speed is typically less than the headline or advertised speed. In Singapore, SamKnows and the IDA (the Singaporean communications regulator) have found in June 2012 that average

broadband speeds typically met advertised rates for in-country traffic. However, international traffic falls significantly below advertised speeds.

In the UK, SamKnows and Ofcom (the UK communications regulator) has found in May 2012 that average xDSL speeds are significantly lower than advertised. However, cable and FTTC services largely meet advertised speeds.

As a consequence of these respective studies the relevant governments have been able to work with industry, ISPs, academics and consumer groups to not only educate the various stakeholders about the limitations of the various technologies available (DSL, Cable, Fibre, Satellite, etc), but also work to promote investment in faster and more consistent consumer broadband generally.

It is the hope of SamKnows that this first report will fill an information gap and provide a reliable reference point and data set for the continued study of internet performance and benchmarking, openness and transparency. As in all its projects SamKnows maintains an open methodology and welcomes the participation of all stakeholders: industry, ISPs, academics, governments and consumer groups. This document details both the results of the study (the data) and the methodology that has been used to collect, aggregate and present the data. The format of this report has been developed over a number of similar studies undertaken by SamKnows since 2009 and will be used as a framework for similar reports and publications over the next three years.

About SamKnows

Since 2009, SamKnows has been developing a comprehensive, transparent and open test methodology which includes every aspect of measuring consumer broadband performance, including: recruiting a panel of consumer volunteers, the tests which are run on specially configured hardware devices on consumer internet connections, the mechanism for collecting and aggregating the data and finally the format for presenting the data and every aspect of this methodology has been scrutinized by project stakeholders representing industry, ISPs, academics, governments and consumer groups from around the world. Importantly, the same tests, hardware and collection methodology are used in all of SamKnows various projects around the world and form a standard test suite and platform.

This European Commission study required for a panel of approximately 10,000 consumers across 30 countries, which include:

- | | | |
|-------------------|----------------|--------------------|
| 1. Austria | 11. Germany | 21. Netherlands |
| 2. Belgium | 12. Greece | 22. Norway |
| 3. Bulgaria | 13. Hungary | 23. Poland |
| 4. Croatia | 14. Iceland | 24. Portugal |
| 5. Cyprus | 15. Ireland | 25. Romania |
| 6. Czech Republic | 16. Italy | 26. Slovakia |
| 7. Denmark | 17. Latvia | 27. Slovenia |
| 8. Estonia | 18. Lithuania | 28. Spain |
| 9. Finland | 19. Luxembourg | 29. Sweden |
| 10. France | 20. Malta | 30. United Kingdom |

Across the European Union SamKnows carries out measurements on xDSL, Cable and FTTx access technologies. These saw average peak download speeds of 7.19Mbps, 33.08Mbps and 41.02Mbps respectively, although there was much variation between countries. It is important to note that the SamKnows methodology has been designed to ensure the most accurate and independent study of internet performance regardless of access technology and home installation.

Whilst this is the first report in the this SamKnows/European Commission study, more are planned during the course of the three year study and SamKnows invites the participation of all stakeholders, for consumers to volunteer, ISPs to review the data and assist in promoting the project, academics to review the methodology and consumer groups to advise on how best to present the findings to European consumers.

Currently the study includes over 250 ISPs, but it is anticipated that more will be able to participate as the project develops.

Overview of Methodology

The purpose of the study is to measure actual consumer experience of effective throughput speed of a representative sample of fixed broadband consumers in all of the 27 Member States of the European Union, as well as: Croatia, Iceland and Norway.

The study uses specially configured SamKnows Whiteboxes (hardware monitoring units), which are placed in the homes of consumer volunteers, selected as part of a representative sample of European fixed broadband

consumers. The purpose of the study is to compare the actual achieved speed, both in rural and urban areas of each country and between countries. The study focuses on wired (fixed) technologies, specifically xDSL, Cable and FTTx and excludes wireless (mobile). The ISPs studied are those considered to have significant market power on their national markets, limited to the two or three largest ISPs where relevant.

The consumers chosen for the panel have been selected on the basis of their geographical balance and whether they fit a pre-defined sample plan, designed to enable the comparison of effective broadband speeds at both a national level and urban/rural areas. It is not the purpose of the study (currently) to compare actual speeds between ISPs within (or between) the countries.

As set out in section B.3, there is not a minimum number of consumers required per country, however the composition of the sample in each Member State should ensure a proper representativeness of the broadband consumers in that country. Each stage of the sample plan design and recruitment of the panel was carried out to ensure the robustness of the results, this includes: the recruitment methodology, breakdown by age, gender, working status, geographical location, ISP, service tier, etc, as well as the specific products or service tiers to be studied.

In all cases, the sample plan has been prepared by SamKnows and approved by the European Commission and the final choice of service tiers to be measured and volunteer consumers has been made on the basis of those which are regarded as the most popular in each country. At the time of creating the sample plan, two thirds of fixed broadband lines in the EU offered speeds of between 2 and 10 Mbps.

At all stages the confidentiality and privacy of the participant consumers was respected and protected. The only data collected by the Whiteboxes was that generated by the tests. SamKnows does not collect any personal data other than that required for the successful completion of the study. SamKnows does not monitor the user's internet traffic.

Once installed, the SamKnows Whitebox runs its tests according to the European Commission test schedule, twenty four hours per day and seven days a week. This enables the analysis of how performance varies according to time of day and day of the week. Key to this is the ability to compare on and off peak performance which is often the most important measure for consumers who are as interested in how consistent their internet connection is (during peak) as they are how fast it is (during off peak). The same test suite is used across all countries, which means that all ISPs are tested identically. The test conditions and frequency are comparable within each country and across all countries.

In the recruitment phase SamKnows uses its own web-based speed test to ensure that each volunteer is consistent with the requirements of the sample plan. This test also collects the IP address of the broadband connection which allows the analysts at SamKnows to identify both the ISP and an approximation of the service tier.

In terms of tests, SamKnows monitors the following indicators in this study:

- Web browsing
- Voice over IP
- Download speed
- Upload speed
- UDP latency
- UDP packet loss
- DNS resolution

In terms of results presentation, this report allows for benchmarking of actual broadband performance across the European Union. The report is not intended to be used by EU consumers to compare the performance of different ISPs, rather for the purposes outlined above.

The SamKnows solution has been designed solely from the point of view of accurately measuring actual customer experience 24x7x365, and presenting this data in a way that makes it most relevant.

It is acknowledged, that further to feedback from local Regulatory Authorities, that enhancements can be made to the sample plan and data processing. To achieve this, a Steering Committee has been proposed, the purpose of which will be to facilitate recruitment and the availability of data that is not currently in the public domain, such as subscriber numbers and distribution, to enable more effective weighting of the data, for example.

Key Findings

Summary

This study presents the results of measurements taken from 9,104 measurement devices in March 2012. These devices were spread across 30 states, distributed according to the sample plan methodology discussed above.

Unless otherwise stated, figures in this report refer to performance at peak times, which is defined as 7pm to 11pm (local time).

The average download speed across all countries was 19.47 Mbps during peak hours, and this increased slightly to 20.12 Mbps when all hours were considered. This figure represents 74% of the advertised headline speed. Note that these are the overall results of the sample, and do not refer to the actual composition of the broadband market across each country

Whilst this compares poorly to the USA's average of 96% of advertised download speed, it is imperative to note that the actual download speeds attained in Europe were considerably higher than those in the USA. xDSL services averaged 7.20Mbps in Europe and 5.30Mbps in the US. Cable services averaged 33.10Mbps in Europe and 17.00Mbps in the US. The same pattern was found for FTTx services too, with Europe averaging 41.02Mbps and US achieving 30.20Mbps.

The scenario is very different when we consider advertised upload speeds. Broadband services are commonly sold with asymmetric download and upload speeds, with the upload speeds being far lower than the download speeds. Across Europe, the average upload speed was 6.20 Mbps, representing 88% of advertised upload speeds.

The actual upload speeds varied very significantly between countries, with those having a large FTTx footprint seeing far higher results than those that do not. However, the performance as a percentage of advertised was broadly similar across all countries and technologies, with all achieving 75% or higher.

The remaining metrics covered in this study provide other indicators of broadband performance. It should be noted that ISPs do not typically advertise expected levels of performance for these metrics, so it is impossible to compare actual versus advertised levels for these. Because of this, the study presents the figures directly and discussion of what certain levels will mean to consumers.

Performance by Technology

The headline figure of 74% of advertised download speed masks many interesting underlying observations.

Firstly, there is significant variation in the performance of different technologies. xDSL based services achieved 63.3% of the headline download speed, whilst cable and FTTx services achieved 91.4% and 84.4% respectively. Note that FTTx includes not only FTTH but also VDSL.

All technologies saw a small decrease of 2-3% during peak hours. The uniformity of this demonstrates that no one technology was more susceptible to peak time congestion than another, and this figure was driven purely by how ISPs engineer their networks.

FTTx services achieved the fastest speeds in absolute terms, at 41.02Mbps. Cable services achieved 33.08Mbps, whilst xDSL services lagged far behind at 7.19Mbps on average.

Cumulative distribution charts presented later in the report demonstrate that the vast majority of cable and FTTx users saw similar speeds to one another. However, xDSL services delivered a very wide variation in download speed to users. This is expected, as speed over xDSL is largely a function of the length of the copper phone line.

Upload speeds varied significantly by technology. FTTx services achieved the highest speeds by far, at 19.8Mbps. This is caused by the fact that many FTTx services across Europe are symmetric (providing the same download and upload speed), or at least provide an upload speed far closer to the download speed. Cable and xDSL services achieved a modest 3.68Mbps and 0.69Mbps in comparison.

Metrics such as latency and packet loss should not be overlooked, as these are just as important (if not more so) to many online activities as download speed is. The average latency across Europe was 33.11ms. This figure is largely dictated by the technology in use, with xDSL averaging 39.94ms and cable and FTTx averaging 24.87ms and 22.02ms respectively. The higher latency for xDSL is expected; the length of the copper phone line coupled with techniques for increasing line stability and reducing packet loss (such as interleaving) are common causes.

Packet loss was found to be 0.5% across Europe, and within this the figure was considerably higher for xDSL services (0.7%) than cable and FTTx (both 0.3%). Interestingly, packet loss for xDSL was sharply lower at 0.4% when we consider all hours in the day. This suggests some countries/ISPs saw significant packet loss at peak times, which drove up the European average.

Section C provides a detailed analysis of performance across Europe as a whole, split by the technologies in use.

A.2.3 **Performance by State**

Download performance varied considerably by state. 16 of the 30 states sampled achieved at least 80% of the advertised download speed. However, two states failed to achieve 50% of the advertised download speed.

This situation is primarily driven by the technologies that have historically been deployed in those countries. The two countries that did not achieve 50% were vastly dominated by xDSL services and – critically – advertised their services using only a handful of very high headline speeds. In France for example, the vast

majority of xDSL services were being advertised with a headline speed of 24Mbps or 28Mbps.

Ofcom has previously carried out studies in the United Kingdom in which their data was weighted for the xDSL operators to normalise for distance from the exchange. Weighting by line length specifically addressed differences in performance that could have been introduced as a consequence of line length. This was not carried out as part of this study, hence some differences in results between these reports.

Whilst it may be tempting to assume that the use of xDSL services automatically mean consumers will not get the speeds they pay for, this is not universally the case. Hungary, which achieved 94% of advertised speed, provides a good counterexample to this: xDSL services achieved over 90% of advertised speeds in this country. This suggests that ISPs in this country were more conservative with their marketing of products.

Performance across other metrics varied considerably by country, with a large proportion of that variation being driven by the technology in use. This is not always the case though; such as with Italy's increase in packet loss from 0.81% to 1.76% during peak hours. Such changes are not driven by technology differences alone.

Section D provides a comprehensive comparison of performance between different states.

A.3 **Next Steps**

This document represents the first in a series of three studies into the performance of broadband services across Europe. The next report is due in the autumn of 2013. In advance of the report, BEREC and the European Commission have agreed to set up a group of experts to participate in the recruitment of consumers and validate data collected, in order to make the publication as accurate and comprehensive as possible.

Following the publication of this first report we anticipate receiving considerable feedback from multiple stakeholders, including national regulators from member states, ISPs, content providers, academics and the public. Working with the European Commission, we will review this feedback and incorporate any changes that are appropriate.

Additionally, SamKnows will also be soliciting for new volunteers to ensure that the customer panel continues to be representative of European broadband services and has sufficient sample sizes. This may require additional targeted deployment of measurement devices where new services are being deployed and become popular.

SamKnows hopes to start building a collaborative working group with ISPs from across the 30 states. This would allow for the methodology to be better targeted

to individual countries' and ISPs' environments. It may also provide a platform for SamKnows to more easily recruit future volunteers to participate in the project. This model has proven very successful in SamKnows' similar project in the USA with the FCC.

Finally, SamKnows plans to expand our footprint of measurement servers, including the deployment of measurement servers within the ISPs' networks. Measurement devices would then be configured to run tests to a set of servers outside of the control of ISPs ("off-net") and also to a set of servers inside their ISP's network ("on-net"), if present. This approach has worked very well in other projects, as it provides a mechanism to health-check the measurement servers themselves and verify any performance anomalies observed.

B

Methodology and Definitions

B.1

Methodology

B.1.1

An Open and Transparent Test Methodology

One of the founding principles of SamKnows is a commitment to open data and a transparent technical methodology. SamKnows is working with academics, governments, industry and consumers worldwide to design and build standard test methodologies and open datasets.

Key to this is releasing all of the technical methodology used to create the data, including making the tests available via open source. It is imperative that the SamKnows tests should be replicable so that the data can be independently verified. This means that rather than operate a closed platform which only uses proprietary code, SamKnows actively looks to publish as much information about its working practices as possible, including making our source code available for independent review.

SamKnows has also introduced a number of documents and processes to ensure that all of our projects are run in a way which is compliant with the SamKnows principles of openness and transparency. For example, in partnership with the FCC and the leading American ISPs, SamKnows introduced a code of conduct to ensure that all participants in the Measuring Broadband America project act in good faith in support of the overall goals of the program. It is hoped that similar documentation can be introduced to the European Commission study.

B.1.2 Measurement Methodology

This section describes the system architecture and network programming features of the tests, and other technical aspects of the methods employed to measure broadband performance during this study.

Hardware vs. Software

A fundamental choice when developing a solution to measure broadband performance is whether to use a hardware or software approach.

Software approaches are by far the most common and allow a very large sample to be reached relatively easily, web-based speed tests fall into this category. These typically use Flash or Java applets, which execute within the context of the user's web browser. When initiated, these clients download content from remote web servers and measure the throughput of the transfer. Some web-based speed tests also perform upload tests, while others perform basic latency checks.

Other less common software-based approaches to performance measurement involve installing applications on the user's workstation, which periodically run tests while the computer is switched on.

All software solutions implemented on a consumer's computer, smart phone, or other internet access device suffer from the following disadvantages for the purposes of this study:

- The software typically does not account for multiple machines on the same network;
- The software may be affected by the quality and build of machine;
- Potential bottlenecks (such as wireless equipment, misconfigured networks, and older computers) are generally not accounted for and result in unreliable data;
- A consumer may move the computer or laptop to a different location which can affect performance;
- The tests may only run when the computer is actually on, limiting the ability to provide a 24-hour profile;

For manually-performed software tests, panelists may introduce a bias by when they choose to run the tests (e.g., may only run when they are encountering problems with their service).

In contrast, hardware approaches involve placing a device inside the user's home that is physically connected to the consumer's internet connection, and periodically running tests to remote targets on the internet. These hardware devices are not reliant on the user's workstation being switched on, and so allow results to be gathered throughout the day and night. The primary disadvantages of a hardware approach are that this solution is much more expensive than a software approach and requires installation of the hardware by the consumer or a third party.

Also, some ISPs in Europe supply customers with a combined modem/router with integrated IPTV support via a dedicated port. In such cases, it is not always possible for the customer to disconnect their TV set-top box from this port and reconnect it to the Whitebox as per our installation instructions. For this reason, it is possible that a customer could be watching IPTV whilst the Whitebox is running tests, which could distort some test results. However the key to understanding the impact of IPTV is an ability to profile the performance of an IPTV-enabled internet connection. It is then possible to spot for performance variation that is as a consequence of IPTV, rather than network congestion. This is something that is being developed by SamKnows analysts, with the intention of this functionality being built-in to the user reporting.

Key features

The SamKnows Performance monitoring framework is a distributed network of Whiteboxes in actual consumers' homes, and is used to accurately measure the performance of fixed line broadband connections based on real-world usage. These are controlled by a cluster of servers, which host the test scheduler and the reporting database. The data is collated on the reporting platform and accessed via a reporting interface and secure FTP. The framework also includes a series of speed-test servers, which the nodes call upon according to the test schedule.

The following technologies are used: Linux, C, Shell scripting, Apache, PHP 5, MySQL, Ajax.

Technical framework

The SamKnows framework solution has been developed since 2008, and currently includes the following 20-point checklist:

SamKnows Technical Objectives	SamKnows Solution
1. Must not change during the Monitoring Period.	The pulling data process is automatic and consistent throughout the monitoring period.
2. Must be accurate and reliable.	Based on independent testing, the hardware solution is reliable.
3. Must not interrupt or unduly degrade the consumer's use of their broadband connection.	The volume of data does not interfere with the broadband experience as tests are not run when a panelist is using their connection.
4. Must not allow collected data to be distorted by any use of the broadband connection by other applications on the host PC and other devices in the home.	The hardware solution does not interfere with the PC and is not dependent on PC. Its only dependence is that the router needs to be switched on as well as the Whitebox.
5. Must not rely on the knowledge, skills and participation of the consumer for its ongoing operation once installed.	The Whitebox is "plug-and-play".
6. Must not collect data that might be deemed personal to the consumer without their consent.	The consent of the consumer regarding the use of their personal data as required by relevant legislation.
7. Must be easy for a consumer to completely remove any hardware and/or software components of the solution if they do not wish to continue with the research programme.	The hardware solution can be disconnected at any time from the home router. As soon as the router is reconnected the connection is resumed as before.
8. Must be compatible with a wide range of xDSL and DOCSIS modems.	The hardware solution can be connected to any router with Ethernet ports.
9. Where applicable, must be compatible with a range of computer operating systems, including but not limited to, Windows XP, Windows Vista, Windows 7, Mac OS and Linux.	The hardware solution is independent of PC operating system and therefore includes all current market standards.
10. Must not expose the consumer's PC to increased security risk, i.e., it should not be susceptible to viruses, it should not degrade the effectiveness of the user's existing firewalls, anti virus and spyware software etc.	Most user firewalls, antivirus and spyware systems are PC based. The hardware solution is plugged in before the PC. Its activity is transparent and does not interfere with those protections.
11. Must be upgradeable from the remote control centre if it contains any software or firmware components.	The Whiteboxes are controlled centrally for updates without involvement of the consumer PC, providing the Whitebox is switched on and connected.

12. Must be removable from the remote control centre if it is a software only solution.	N/A, the Whitebox is hardware- based.
13. Must identify when a user changes broadband provider or package (e.g. by a reverse look up of the consumer's IP address to check provider, and by capturing changes in modem connection speed to identify changes in package).	Regular monitoring of any changes in speed, ISP, IP address or performance. Should a consumer change package, they will be invited to notify us of the change or confirm that no change took place since the last report.
14. Must permit, in the event of a merger between ISPs, separate analysis of the customers of each of the merged ISP's predecessors.	Data is stored based on the ISP of the panelist, and can therefore be analyzed individually or as a whole.
15. Must identify if the consumer's computer is being used on a number of different networks (e.g., if it's a laptop).	The Whitebox is not PC or laptop dependent, but is broadband connection dependent.
16. Must identify when a specific household stops providing data.	The Whitebox needs to be connected and switched on to pull data. If it is switched off or disconnected its absence is detected at the next data pull process.
17. Must not require an amount of data to be downloaded which may materially impact on any data caps or fair usage policy the ISP has imposed on the end user, or trigger traffic shaping policies to be implemented by the ISP.	The data volume generated by the information collected does not exceed any policies set by ISPs. Panelists with bandwidth restrictions can have their tests set accordingly.
18. Must ensure that its tests are run in a manner which does not make it possible for ISPs to identify the broadband connections which form their Panel and therefore potentially enable ISPs to "game" the data by providing a different quality of service to the Panel members and the wider customer base.	The data packet profile is not identifiable unless it is subject to a DPI process that specifically looks for these profiles. This can only be done if the ISPs are aware of the profile of the data and if the ISP has a level of resources sufficient to monitor its entire customer base.
19. Must be consistent and adhere to all relevant standards for internet measurement.	The measurement platform is being used as the basis for the development of global standards.
20. The solution must be sufficiently scalable to become a global measurement platform.	The performance measurement platform has been designed to be a global platform that can scale to many millions of customers.

B.1.2.1 **Fixed broadband hardware probes**

SamKnows uses hardware probes (Whiteboxes) for the purpose of accurately measuring end-user broadband performance. For this study, there are two types of probes, subject to the achievable speed of the internet connection.

The Whiteboxes execute a series of software tests over the broadband connection they are connected to. The results of these tests are reported securely up to hosted backend infrastructure.

The majority of tests run against a network of test nodes. These are dedicated servers either “on-net” (on the local ISP’s network) or “off-net” (on the public internet). Some tests will execute against real applications hosted on the internet, mimicking their behaviour and measuring key performance variables.

When a testing cycle has completed, the results are encrypted and transmitted over SSL to hosted backend infrastructure for processing and presentation through a web interface to each panellist and other interested parties.

Panellists are, as part of the terms of service, required to leave their Whitebox and other networking equipment permanently powered on and connected to ensure consistent testing.

All SamKnows Whiteboxes run a custom distribution of Linux, derived from OpenWrt. Many standard OpenWrt features have been removed to save space on the device, and some additional features have been added to support the measurements.

The custom firmware is flashed at the factory and is not directly upgradeable by the user hosting the Whitebox. The firmware is remotely upgradeable by SamKnows.

This cut-down operating system provides network connectivity and the measurement applications alone – there is no web interface and the Whitebox provides no routing functionality. Panellists have no ability to disable, reconfigure or influence the SamKnows software in any way through normal usage.

SamKnows’ firmware makes use of GPL v2.0 licenced code. The source code for SamKnows’ firmware build is available at: <https://files.samknows.com/~gpl/>

All communications between the Whitebox and the Data Collection Service on the backend hosted infrastructure are initiated by the Whitebox, encrypted over SSL and subject to authentication

The Whitebox communicates with the target test nodes over a variety of TCP and UDP ports. The Whitebox will also communicate with some unmanaged services over both TCP and UDP.

The SamKnows software suite has the ability to auto-update itself, downloading updated binaries and testing schedules from the Data Collection Service and storing locally in RAM or flash.

Whitebox 1.0 (SK-TL-WR741ND)

The SK-TL-WR741ND can accurately measure fixed-line broadband connections of up to 100Mb/s. Like all SamKnows Whiteboxes, it operates Linux using a 2.6.x kernel. The specifications of the device are as follows:

- 5x 100Mbps Ethernet
- 1x 802.11n wireless interface
- Single DC power (9V @ 750mA)
- Dimensions: 174mm x 118mm x 33mm
- Power draw: 4W
- Weight: 500g


Whitebox 2.0 (SK-TL-WR1043ND)

The SK-TL-WR1043ND can accurately measure fixed-line broadband connections of up to 250Mb/s. Like all SamKnows Whiteboxes, it operates Linux using a 2.6.x kernel. The specifications of the device are as follows:

- 400Mhz MIPS CPU
- 32MB RAM
- 5x 1Gbps Ethernet
- 1x 802.11n wireless interface, 3 antennas
- Single DC power (12V @ 1500mA)
- Power draw: 5W
- Weight: 512g

Installation

The Whitebox operates as an Ethernet bridge, co-existing with an existing modem/router. All wired devices should connect through the Whitebox. Wireless devices should continue to connect to their existing router:



Above: The devices are installed on a consumer's broadband connection (which will likely be in use through normal day-to-day activity). Note that it is necessary for the Whitebox to take some precautions in order to protect the validity of the gathered data. In order to determine when it is safe to execute tests, the end user's traffic levels are monitored continuously. This is the reason for connecting all Ethernet devices through the Whitebox.

In the scenario above, the Whitebox will passively monitor the strongest wireless access point in the vicinity for traffic. It passively listens for packets on that network, and if packets are flowing, then measurements will not be executed.

B.1.2.2 Software test (to initially check volunteer information)

There are some situations where this level of certainty over the results is less critical and the focus is instead upon providing instantaneous data to end users. SamKnows use a web-based test as a key stage in the recruitment of consumer volunteers. As part of the multi-step recruitment process each applicant is asked to complete a web-based speed test, the results of which are then compared against the European Commission sample plan to determine whether the volunteer matches the sample plan criteria.

To address this requirement, SamKnows has developed a Java-based software application for measuring selected broadband performance metrics. The measurements have been built to the same specification as those used in the SamKnows Whiteboxes, but are designed to run in software on a user's workstation.

The Java-based application is typically embedded inside web pages as an applet, providing widespread platform support (Windows, Mac OS, Linux).

The following key metrics are included in the software-based measurements:

Metric	Primary measure(s)
Download speed	Throughput in Megabits per second utilising one or more concurrent TCP connections
Upload speed	Throughput in Megabits per second utilising one or more concurrent TCP connections
UDP latency	Average round trip time of a series of randomly transmitted UDP packets
UDP packet loss	Percentage of UDP packets lost from latency test

The Architecture

The web-based test is embedded as a Java applet in a publicly accessible web page. In its simplest use case, end users will initiate the measurement process by clicking on the ‘Start’ button. A short while later the measurement results will be reported to the end user.

Determining the best measurement server

Upon start up, the application runs a brief latency measurement to all measurement servers hosted by SamKnows. This process allows us to determine the nearest measurement server (in terms of latency) The measurement server with the lowest round-trip latency is selected as the target for all subsequent measurements (throughput, latency and packet loss).

Additionally, if the ISP has installed “on-net” measurement servers within their network then the application will also select the nearest one of these servers. Measurements are run against both the “on-net” and off-net servers.

Cross-traffic, in-home network issues and configuration differences

One of the key advantages of the hardware-based Whitebox is its ability to detect cross-traffic and defer tests. Furthermore, its position within the home network (connected directly to the modem or gateway) means that it is unaffected by in-home network issues (such as those caused by wireless networks).

A purely software-based approach is not able to account for such issues. However, we can apply a number of mechanisms in an attempt to reduce or detect their impact.

Cross-traffic within the local client (e.g. PC) is measured and tests will not be executed if the client is transferring more than 64kbit/s.

Additionally, the web-based test will poll the user’s gateway via UPnP for traffic counters. This allows for cross-traffic within the home to be fully accounted for, and measurements will not be executed if the gateway is transferring more than 64kbit/s. However, this UPnP-based approach is far from universally supported. A recent study (February 2012) showed that approximately 22% of gateways in

Europe supported traffic counter reporting by UPnP, but this figure is expected to rise.

In-home network issues (such as poor wireless) cannot be excluded by the web-based test. However, we can attempt to identify them. In particular, the web-based test records the connection media used by the client and its connected speed (e.g. Ethernet at 100Mbps, or Wireless at 54Mbps). Additionally, the web-based test will also run a brief ICMP latency and packet loss measurement to the user's gateway. If this reports more than 2ms latency and 0% packet loss, then the measurements are aborted with a message stating that the user's in-home network appears to be operating poorly.

Client configuration issues (such as insufficient TCP settings, firewall products, RAM or CPU) are checked for before measurements begin. If these fall outside of accepted bounds then the tests are aborted and the user is informed.

In all of the error conditions above the user will be informed of the reason why the measurements were not executed. The user may override the failure and run the measurements anyway, but the results will be recorded on the server side with a 'tainted' flag indicating that they were not run under optimal conditions.

Capturing location and ISP data

The approximate location of the client is determined through two means:

Firstly, the server side examines the IP address of the client and utilises geo-location databases such as Maxmind to find the location and ISP of the user. The physical location is typically accurate to city-level and the ISP can be determined with near 100% accuracy.

Additionally, if the client PC has 802.11 wireless support then the list of nearby wireless peers points is used in conjunction with an online service to determine a more accurate physical location. This typically provides accuracy to street or postcode level.

Communications

All communications between the web-based test and the Data Collection Service on the backend hosted infrastructure are initiated by the software application and encrypted over SSL.

The software application communicates with the measurement servers over a variety of TCP and UDP ports. ICMP is also used to determine the server with the lowest round-trip latency.

B.1.2.3 Overview of network test nodes

Whiteboxes target dedicated, bare metal servers configured as the end point for the speed, streaming, VoIP, jitter, latency, packet loss and availability tests.

Whiteboxes query the backend infrastructure to find out which target test node to test against, so the test nodes targeted can be fully managed and updated dynamically.

SamKnows has been a member of the Measurement Lab research consortium (M-Lab) since 2009. Alongside SamKnows dedicated test nodes, we can also use the M-Lab infrastructure as destinations for our remote tests during this project. These nodes are located in ten major global internet peering locations.

An important aspect of the SamKnows methodology is that both ends of the test are controlled by the SamKnows measurement platform. In fact, the server-side component of the test is as important as the client-side software. Each network test node is built with a standard specification and loaded with proprietary SamKnows system packages.

The Whiteboxes target these dedicated, bare metal servers configured as the end point for the speed, streaming, VoIP, jitter, latency, packet loss and availability tests.

On-network and off-network test nodes

SamKnows maintains a global network of test nodes that the Whiteboxes test against. Many of these are built upon the Measurement Labs infrastructure and their locations can be found at <http://code.google.com/p/ndt/wiki/MLabOperations>. These nodes are said to be “off-net”, as they do not reside directly on any one ISP’s network.

Please note that all tests run against M-Lab test nodes are subject to M-Lab’s data release policy, which requires publication of all data collected within 1 year of collection.

ISPs may contribute hardware for the purposes of hosting “on-net” test nodes. These are nodes which are hosted within the ISP’s network. The purpose of these nodes is to allow the ISP to determine what (if any) degradation in performance occurs outside of their network.

This first European Commission study incorporates only off-net test nodes.

At start up, Whiteboxes retrieve a list of all active test nodes from the SamKnows infrastructure. The Whitebox then uses a simple series of ICMP pings to measure approximate latency to each. The node with the lowest latency is said to be the “closest” and will be used from that point on. Whiteboxes will then perform tests against the closest off-net node and the closest “on-net” node for that ISP (assuming the ISP has provided one). Should the selected test node become unavailable for an extended period then SamKnows will ask the Whitebox to re-select its closest targets.

In the European Commission study ISPs have not yet contributed “on-net” measurement servers. All results presented here are taken from “off-net” measurement servers. SamKnows encourages ISPs to consider providing “on-

net” servers for future studies in order to add additional checks and balances to the measurement process.

Test node specification

Test nodes must meet the following minimum specification:

- Dual-core CPU of 2Ghz
- 4GB RAM
- 80GB disk space
- Gigabit Ethernet connectivity, with gigabit upstream link
- Centos/RHEL 5.x/6.x

B.1.2.4 European Commission Study Test Nodes

Measurement servers geographically distributed across Europe were used to conduct the client to server tests. These measurement servers were located at major European peering and Internet exchange points (IXPs). The measurement clients always chose the nearest server (in terms of round-trip latency) to execute their tests against. Ensuring the measurement server is nearby helps keep the number of intermediate networks low, thus reducing the chance that measurements will be negatively influenced by a congested upstream network.

Measurement servers were located as follows:

Location	Hostnames
Amsterdam, Netherlands	ispmon.samknows.mlab1.ams01.measurement-lab.org ispmon.samknows.mlab2.ams01.measurement-lab.org ispmon.samknows.mlab3.ams01.measurement-lab.org n1-amsterdam-nl.samknows.com
Athens, Greece	ispmon.samknows.mlab1.ath01.measurement-lab.org ispmon.samknows.mlab2.ath01.measurement-lab.org ispmon.samknows.mlab3.ath01.measurement-lab.org
Hamburg, Germany	ispmon.samknows.mlab1.ham01.measurement-lab.org ispmon.samknows.mlab2.ham01.measurement-lab.org ispmon.samknows.mlab3.ham01.measurement-lab.org
London, UK	ispmon.samknows.mlab1.lhr01.measurement-lab.org ispmon.samknows.mlab2.lhr01.measurement-lab.org ispmon.samknows.mlab3.lhr01.measurement-lab.org n1-the1.samknows.com n2-the1.samknows.com
Madrid, Spain	ispmon.samknows.mlab1.mad01.measurement-lab.org ispmon.samknows.mlab2.mad01.measurement-lab.org ispmon.samknows.mlab3.mad01.measurement-lab.org n1-madrid-es.samknows.com
Milan, Italy	ispmon.samknows.mlab1.mil01.measurement-lab.org ispmon.samknows.mlab2.mil01.measurement-lab.org ispmon.samknows.mlab3.mil01.measurement-lab.org
Turin, Italy	ispmon.samknows.mlab1.trn01.measurement-lab.org ispmon.samknows.mlab2.trn01.measurement-lab.org ispmon.samknows.mlab3.trn01.measurement-lab.org
Paris, France	ispmon.samknows.mlab1.par01.measurement-lab.org ispmon.samknows.mlab2.par01.measurement-lab.org ispmon.samknows.mlab3.par01.measurement-lab.org n1-paris-fr.samknows.com
Bucharest, Romania	n1-bucharest-ro.samknows.com

Gdansk, Poland	n1-gdansk-pl.samknows.com
Warsaw, Poland	n1-warsaw-pl.samknows.com
Zagreb, Croatia	n1-unizg-zagreb-hr.samknows.com
Hudiksvall, Sweden	n1-hudiksvall-se.samknows.com
Riga, Latvia	n1-riga-lv.samknows.com


B.1.2.5 Control Suite

SamKnows provides a fully managed solution. All areas of the measurement platform are remotely configurable and the Whiteboxes are managed centrally via the SamKnows Control Suite. This allows the Administrator to remotely amend the test schedule, remove tests or introduce new tests.

- The Whiteboxes are shipped without test binaries; instead they contain an executable allowing them to bootstrap remotely.
- Once the Whitebox is installed for the first time by the panellist, it automatically communicates with the reporting platform, which authenticates the Whitebox and remotely installs the latest version of the test software.
- Periodically, each Whitebox re-checks with the reporting platform, via SSL and, if appropriate, downloads the latest version of the test software and schedule.
- During these periodic checks, the Whitebox uploads data from the previous period, which is then pushed to the Reporting Engine.

An administrator can perform a number of actions using the Control Suite including:

- Add new Whitebox to the system
- Assign users and reporting engine logins
- Update package information
- Update unit assignments
- View data across all panellists in aggregated form



A comprehensive measurement suite

The SamKnows methodology and platform has been designed to be flexible enough to allow for whatever future modifications or enhancements are required, but at the same time the ‘out of the box’ solution provides a fully inclusive package of every available performance measurement test. The table below details the tests included in the study:

Metric	Primary measure(s)
Web browsing	The total time taken to fetch a page and all of its resources from a popular website
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
Download speed	Throughput in Megabits per second utilising three concurrent TCP (Transmission Control Protocol) connections
Upload speed	Throughput in Megabits per second utilising three concurrent TCP connections
UDP (User Datagram Protocol) latency	Average round trip time of a series of randomly transmitted UDP packets
UDP packet loss	Percentage of UDP packets lost from latency test
DNS (Domain Name Server) resolution	The time taken for the ISP's recursive DNS resolver to return an A record for a popular website domain name

Testing schedule

A test cycle on the Whitebox occurs once an hour every hour, 24 hours a day. The timing of the testing is randomised per Whitebox to ensure an even spread across the panel.

A scheduling service on the Whitebox manages the following tasks:

- Execute tests
- Send Test results
- Check the backend service for a new testing schedule
- Check the backend service for updated performance tests

The availability and data usage tests run permanently in the background as soon as the Whitebox has booted.

A test cycle may last up to 15 minutes depending on the performance of the host internet connection. Once testing is complete, results are securely transmitted to the Data Collection Service on the backend infrastructure.

The following schedule provides a breakdown of test durations and indicative impact on monthly bandwidth usage.

Testing Schedule for Europe (Fixed)

Test Name	Test Target(s)	Test Frequency	Test Duration	Est. Daily Volume
Web browsing	3 popular websites	Hourly, 24x7	Est. 3 seconds	8.4MB
Voice over IP	1 off-net test node	Every other hour, 24x7	Fixed 10 seconds at 64k	1.92MB
Download speed	1 off-net test node	Every other hour, 24x7	9MB fixed size	108MB
Upload speed	1 off-net test node	Once 12am-6pm Once 6am-12pm Once 12pm-6pm Once 6pm-12pm	3MB fixed size	12MB
UDP latency	1 off-net test node	Hourly, 24x7	Permanent	1MB
UDP packet loss	1 off-net test node	Hourly, 24x7	Permanent	N/A (uses above)
DNS resolution	3 popular websites	Hourly, 24x7	Est. 1 second	0.1MB
ICMP latency	1 off-net test node	Hourly, 24x7	Est. 5 seconds	0.1MB
ICMP packet loss	1 off-net test node	Hourly, 24x7	N/A (As ICMP latency)	N/A (uses above)

Test Software

SamKnows has designed and developed its software and technology in-house, ensuring adherence to relevant RFCs (Request for Comment). All performance tests are written in C, for performance and portability across a range of hardware platforms.

SamKnows performance tests do not incorporate any third party commercial or free or open source (F/OSS) code. Some tests do however dynamically link to F/OSS libraries.

All times are measured in microseconds.

To provide metrics on the key performance indicators requested, a series of tests are utilised.

Metric	SamKnows Fixed
Web Browsing	✓
Video Streaming	✓
VOIP Emulation	✓
Downstream Throughput	✓
Upstream Throughput	✓
Latency	✓
Packet Loss	✓
DNS Resolution	✓
FTP Throughput	✓
Peer-to-Peer File Sharing	✓
Email Relaying	✓
Latency Under Load	✓
Loss Under Load	✓

Web browsing

The Web browsing test measures the time taken to fetch the HTML and referenced resources from a page of a popular website. This test does not test against centralised testing nodes; instead it tests against real websites, ensuring that content distribution networks and other performance enhancing factors may be taken into account.

Each Whitebox will test three common websites on every test run. The time taken to download the resources, the number of bytes transferred and the calculated rate per second will be recorded. The primary measure for this test is the total time taken to download the HTML page and all associated images, JavaScript and stylesheet resources.

The results include the time taken for DNS resolution. The test uses up to eight concurrent TCP connections to fetch resources from targets. The test pools TCP connections and utilises persistent connections where the remote HTTP server supports them.

The test may optionally run with or without HTTP headers advertising cache support (through the inclusion or exclusion of the “Cache-Control: no-cache” request header). The test is designed to replicate the user experience of Microsoft internet Explorer 8.

Voice over IP

This test utilises the same generic streaming test as the video test, albeit with different configuration. The test operates UDP and, unlike the video streaming test, utilises bi-directional traffic.

The client initiates a UDP stream to the server and a fixed-rate stream is tested bidirectionally. A de-jitter buffer of 25ms is used to reduce the impact of jitter. The test measures this disruption by monitoring throughput, jitter, delay and loss. These metrics are measured by subdividing the stream into blocks, and measuring the time taken to receive each block (as well as the difference between consecutive times).

The test uses a 64kbps stream with the same characteristics and properties (i.e. packet sizes, delays, bitrate) as the G.711 codec.

Jitter is calculated using the PDV (Packet Delay Variation) approach described in section 4.2 of RFC5481. The 99th percentile will be recorded and used in all calculations when deriving the PDV.

UDP latency and packet loss

This test measures the round trip time of small UDP packets between the Whitebox and a target test node. Each packet contains consists of an 8-byte sequence number and an 8-byte timestamp. If a packet is not received back within three seconds of sending, it is treated as lost. The test records the number of packets sent each hour, the average round trip time of these and the total number of packets lost. The test will use the 99th percentile when calculating the summarised minimum, maximum and average results.

The test operates continuously in the background. It is configured to randomly distribute the sending of the echo requests over a fixed interval, reporting the summarised results once the interval has elapsed.

Speed tests

This test measures the download and upload speed of the given connection in bits per second by performing multi-connection GET and POST HTTP requests to a target test node.

Binary non-zero content, herein referred to as the payload, is hosted on a web server on the target test node. The test operates for either a fixed duration (in seconds) or a fixed volume (in MB). It can also output average throughput at multiple intervals during the test (e.g. once every 5 seconds). The client will attempt to download as much of the payload as possible for the duration of the test. The payload and all other testing parameters are configurable and may be subject to change in the future.

Four separate variations of the test are supported:

- Single connection GET
- Multi connection GET
- Single connection POST
- Multi connection POST

Note that SamKnows recommends the usage of the multi connection test for all faster service tiers, and typically uses 3 concurrent connections. Each connection used in the test counts the numbers of bytes of the target payload transferred between two points in time and calculates the speed of each thread as Bytes transferred/Time (seconds).

Factors such as TCP slow start and congestion are taken into account by repeatedly downloading small chunks (default 256KB) of the target payload before the real testing begins. This “warm up” period is said to have been completed when three consecutive chunks were downloaded at the same speed (or within a small tolerance (default 10%) of one another). In a multi connection test, three individual connections are established (each on its own thread) and are confirmed as all having completed the warm up period before timing begins.

Content downloaded is output to /dev/null or equivalent (i.e. it is discarded), while content uploaded is generated and streamed on the fly from /dev/urandom.

The following is an example of the calculation performed for a multi connection test utilising three concurrent connections.

S = Speed (Bytes per second)

B = Bytes (Bytes transferred)

C = Time (Seconds) (between start time point and end time point)

S1 = B1 / C1 (speed for Thread 1 calculation)

S2 = B2 / C2 (speed for Thread 2 calculation)

S3 = B3 / C3 (speed for Thread 3 calculation)

Speed = S1 + S2 + S3

Example values from a 3MB payload:

B1 = 3077360 C1 = 15.583963

$B2 = 2426200$ $C2 = 15.535768$
 $B3 = 2502120$ $C3 = 15.536826$
 $S1 = B1/C1 = 197469.668017$
 $S2 = B2/C2 = 156168.655454$
 $S3 = B3/C3 = 161044.475879$
 $S1 + S2 + S3 = \text{Total Throughput of the line} = 197469.668017 + 156168.655454$
 $+ 161044.475879 = 514682 \text{ (Bps)} * 0.000008 = 4.12 \text{ Mbps}$

DNS resolution

This test measures the DNS resolution time of a selection of common website domain names. These tests will be targeted directly at the ISPs recursive resolvers. A list of appropriate servers will be sought from each ISP in advance of the tests.

Test software summary

As network users, over time, come to expect increasing capability from their networks, the tests used to measure general network speed, multimedia performance and network integrity must be suitably robust and well-designed. All SamKnows tests have been independently certified and/or approved not only by a pool of ISP clients but also by a national internet regulator.

The SamKnows proprietary test suite has been developed to simulate real-world user experience of a broadband service.

SamKnows has designed and developed its performance tests in-house; ensuring adherence to relevant RFCs. Our testing methodology has been independently reviewed by MIT, The Georgia Institute of Technology, ISPs and government regulators.

B.1.2.6 Reporting Infrastructure

SamKnows employs a fully managed infrastructure for the purposes of data collection from the Whiteboxes, data processing, data presentation and Whitebox management.

Currently hosted directly in United Kingdom, the back-end makes use of dedicated hardware firewalls, load balancers and bare metal hardware.

SamKnows operations oversee the management of the backend infrastructure, adhering to industry standard practices for security and operational management.

The backend can be broken down into four distinct areas:

1. **Data Collection Service:** The data collection service, or DCS, is the gateway for the Whitebox to communicate with the back-end for sending test results and requesting configuration updates. Communication with the DCS is over TCP 443 with all communications encrypted via SSL.


2. Data Processing: A cluster of database servers utilising a specialized column based storage engine to process and store results data. All publicly identifiable information (PII) is encrypted and is only accessible by panellists themselves and SamKnows.
3. Data Presentation: Data is made available via a Web 2.0 style-reporting interface, accessible over SSL with granular access controls.
4. Data Feeds: Whilst the dashboard and more graphical data visualisations might be appropriate for the majority of stakeholders, SamKnows also provides feeds via API (Application Programming Interface) to enable more statistical analysis of the data.

Individual consumers: real-time reporting dashboard

Each consumer is able to monitor their own performance data in real-time by logging in securely to their own version of the SamKnows Reporting System.


Performance graphs and statistical results are collected into the areas of speed, integrity and multimedia performance and are viewable both in dashboard overview format and in extended detail in each relevant section.

Each report is fully customizable in the same way as for the Administrator's View, with the limitation that users are only able to view their own data.



Email report card

SamKnows has developed a number of other reporting mechanisms alongside the reporting system. The purpose of these is for ISPs or other service operators to reach out to their customers with performance information. The graphic below illustrates the broadband report card which can be sent out by email automatically:



Smart phone applications

To compliment the web-based reporting system and the monthly report card, SamKnows has also developed a Smartphone App for both iPhone and Android. Participants can use this App to login to view performance data directly from their Smartphone.


01.Login Screen



02.Statistics



03.About



04.Stats/Share



Sample Plan Methodology

This section describes the background to the study and methods employed to design the target panel, select volunteers for participation, and manage the panel to maintain the statistical and operational goals of the program.

To apply our sampling methodology specifically to the EU we extract data from a the European Commission’s study, “Broadband Coverage in Europe” (2010), other European Commission documents such as https://ec.europa.eu/digital-agenda/sites/digital-agenda/files/broadband_country_charts_2011_July.xls and from commercial sources . The data in this study is broken down according to country, technology, speed tier, and rural/urban area. Using this detailed information, we select technology, ISP, and the country as the key primary independent variables. This allows us to measure the direct effect of these variables on broadband speeds. The rural/urban breakdown in the data comprises one of the secondary independent variables that we control for when analysing broadband across the EU.

Within each country and technology type, we then analyse the data taken from the 2010 European Commission study by looking at the percentage of the country’s broadband users within each speed tier. If the percentage of users is less than 5% then this speed tier is excluded from our EU sampling plan. For example, the proportion of broadband users in Belgium with cable technology and a speed tier of 0-2Mbps is 0%, and those with a tier of 2-8Mbps is just 3.27%. Therefore these cells are excluded from our sample. In this way, we apply our sampling methodology to the EU by including only the speed tiers that are important for European users and therefore relevant for measuring broadband speeds, as well as avoiding potential bias resulting from cells with very small numbers of broadband users.

Whilst the 2010 European Commission study that we use to instruct our sampling structure decisions is the most recent and relevant publication, the ever-changing nature of broadband technology combined with the four year time frame of the SamKnows study means that the percentage of users in each part of the breakdown should be expected to evolve over time. In particular, the breakdown of data by speed tier is likely to change, as ISPs inevitably progress towards providing ever-higher speed tiers on the back of technological improvements. To account for this, we continually review, modify and update our sampling structure by gathering data directly from the ISPs themselves relating to the speed tiers they offer their consumers. Note that although the ISPs themselves should be expected to remain fairly stable over time, it is nevertheless likely that the breakdown by technology will change, albeit it much more slowly than speed tiers. We account for this in a similar way — for example, with the current rise of fiber technology deployed to consumers in the UK. In this way, as we apply our sampling methodology and structure to the EU, we are able to continually adjust

and rebalance the structure in line with the current broadband usage, rather than simply rely on a static ‘snapshot’ analysis of broadband based on past data only. SamKnows is therefore able to create a sample as dynamic as the changing nature of the broadband industry itself.

The overall sample plan broken down by country technology and speed tier is shown below:

Note: The sample plan and subsequent examples provided below are based on data collated in 2009 and published in 2010. Because of the ever-changing nature of broadband services, the panel composition will change over time to remain consistent with the overall market. However the sample plan methodology will remain consistent throughout.

		DSL			CABLE				FTTx				TOTAL
		0-2Mbps	2-8Mbps	8-30Mbps	0-2Mbps	2-8Mbps	8-30Mbps	30+Mbps	0-2Mbps	2-8Mbps	8-30Mbps	30+Mbps	
Austria	AT	75	75			75							225
Belgium	BE		75	75			75	75					300
Bulgaria	BG		50	50			50				50	50	250
Croatia	HR		50			50							100
Cyprus	CY	50	50										100
Czech Republic	CZ	75	75			75	75						300
Denmark	DK		50	50		50	50						200
Estonia	EE	50	50			50			50				200
Finland	FI	50	50	50		50							200
France	FR	200	200	200			100	100					800
Germany	DE	200	350	200			200	200					1150
Greece	EL	75	75	75									225
Hungary	HU	75	75			75	75	75					375
Iceland	IS	50	50										100
Ireland	IE	50	50			50							150
Italy	IT	200	350	200									750
Latvia	LV		50	50						50	50		200
Lithuania	LT	50	50				50			50	50		250
Luxembourg	LU		50	50			50						150
Malta	MT		50			50							100
Netherlands	NL	150	150	150		150	150	150					900
Norway	NO	50	50		50						50		200
Poland	PL	250	150			150	150						700
Portugal	PT		50	50		50	50						200
Romania	RO		75			75					75	75	300
Slovakia	SK		50				50				50	50	200
Slovenia	SI	50	50			50					50		200
Spain	ES	200	200	200		200	200						1000
Sweden	SE		75	75		75					75	75	375
United Kingdom	UK		200	200			200						600
TOTAL		1900	2925	1675	50	1275	1525	600	50	100	450	250	10800

Example 1: Belgium

Belgium consists of just over 3 million broadband subscribers, of which 1,760,000 are xDSL subscribers and 1,300,000 are cable subscribers. It is worth noting that xDSL users in Belgium are reported under FTTx because most of the network uses VDSL2. The broadband population is spread as per the table below for each of the 2 technologies per density and per headline speed. As you will see Cable 2Mbps – 8Mbps speed range only consists of 3.27% of the overall population, and therefore is not included in the final sample. The actual composition of the panel in Belgium is shown in B.1.4.

2. BELGIUM					
DSL					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	% of Country
0-2Mbps	0	0	0	0	0.00%
2-8Mbps	520000	600000	50000	1170000	38.24%
8-30Mbps	500000	90000	0	590000	19.28%
TOTAL	1020000	690000	50000	1760000	57.52%
DSL % SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
2-8Mbps	44%	51%	4%	100%	
8-30Mbps	85%	15%	0%	100%	
DSL SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
2-8Mbps	33	38	3	75	
8-30Mbps	64	11	0	75	
CABLE					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	% of Country
0-2Mbps	0	0	0	0	0.00%
2-8Mbps	20000	80000	0	100000	3.27%
8-30Mbps	350000	250000	0	600000	19.61%
30+Mbps	350000	250000	0	600000	19.61%
TOTAL	720000	580000	0	1300000	42.48%
CABLE % SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
8-30Mbps	58%	42%	0%	100%	
30+Mbps	58%	42%	0%	100%	
CABLE SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
8-30Mbps	44	31	0	75	
30+Mbps	44	31	0	75	
PERCENTAGE OF COUNTRY COVERED BY SAMPLE					96.73%
Belgium	Belgacom	Incumbent			
	Telenet	New Entrant			
	Voo	New Entrant			

Example 2: The Netherlands

The Netherlands consists of just over 5.8 million broadband subscribers, of which 3,290,000 are xDSL subscribers and 2,590,000 are cable subscribers. The broadband population is spread as per the table below for each of the 2 technologies per density and per headline speed. As you will see the whole of the countries broadband population is included in the sample. The actual composition of the panel in the Netherlands is shown in B.1.4.

21. NETHERLANDS					
DSL					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	% of Country
0-2Mbps	350000	100000	50000	500000	8.50%
2-8Mbps	1300000	400000	30000	1730000	29.42%
8-30Mbps	750000	300000	10000	1060000	18.03%
TOTAL	2400000	800000	90000	3290000	55.95%
DSL % SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
0-2Mbps	70%	20%	10%	100%	
2-8Mbps	75%	23%	2%	100%	
8-30Mbps	71%	28%	1%	100%	
DSL SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
0-2Mbps	105	30	15	150	
2-8Mbps	113	35	3	150	
8-30Mbps	106	42	1	150	
CABLE					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	% of Country
0-2Mbps	0	0	0	0	0.00%
2-8Mbps	400000	100000	0	500000	8.50%
8-30Mbps	1100000	500000	150000	1750000	29.76%
30+Mbps	340000	0	0	340000	5.78%
TOTAL	1840000	600000	150000	2590000	44.05%
CABLE % SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
2-8Mbps	80%	20%	0%	100%	
8-30Mbps	63%	29%	9%	100%	
30+Mbps	100%	0%	0%	100%	
CABLE SAMPLE BREAKDOWN					
SPEED	URBAN	SUBURBAN	RURAL	TOTAL	
2-8Mbps	120	30	0	150	
8-30Mbps	94	43	13	150	
30+Mbps	150	0	0	150	
PERCENTAGE OF COUNTRY COVERED BY SAMPLE					100.00%
The Netherlands	KPN	Incumbent			
	BBNED / Tele2	New Entrant			
	Online	New Entrant			
	UPC Nederland	New Entrant			
	Ziggo	New Entrant			

Final Panel Composition

As outlined above, the European broadband market is incredibly dynamic with new services frequently being made available. Because of this ever-changing nature of broadband services, and SamKnows constantly looking to track how the market evolves, the final panel composition changed (by design) from the original sample plan. The final panel of valid panelists is therefore displayed below:

Country	xDSL Panel	Cable Panel	FTTx Panel
Austria	41	20	14
Belgium	5	221	110
Bulgaria	17	28	72
Croatia	51	12	0
Cyprus	56	24	0
Czech Republic	114	106	6
Denmark	94	19	50
Estonia	16	5	10
Finland	138	25	35
France	334	26	22
Germany	509	152	89
Greece	209	0	2
Hungary	51	145	10
Iceland	51	0	23
Ireland	89	62	2
Italy	613	0	14
Latvia	42	10	63
Lithuania	53	11	176
Luxembourg	45	12	22
Malta	67	64	0
Netherlands	203	277	91
Norway	64	15	82
Poland	362	173	25
Portugal	88	48	20
Romania	42	36	158
Slovakia	66	35	69
Slovenia	51	21	67
Spain	483	161	72
Sweden	68	30	180
United Kingdom	1080	270	156
TOTAL	5102	2008	1640

B.1.4.1 Use of an All Volunteer Panel

In 2008, SamKnows conducted a test of residential broadband speed and performance in the United Kingdom and during the course of that test determined that attrition rates for such a test were lower when an all-volunteer panel was used, rather than attempting to maintain a panel through an incentive scheme of monthly payments. Consequently, in designing the methodology for this broadband performance study, we relied entirely on volunteer consumer broadband subscribers. The volunteers were selected from a large pool of prospective participants according to a plan designed to generate a representative sample of desired consumer demographics, including geographical location, ISP, and speed tier. As an incentive for participation, volunteers were given access to a personal reporting suite which allowed them to monitor the performance of their broadband service. They were also provided with a measurement device referred to in the study as a “Whitebox,” configured to run custom SamKnows software.

B.1.4.2 Sample Size and Volunteer Selection

The study allowed for a target deployment of up to 10,000 Whiteboxes to volunteer panelists across the European Union. The number of volunteers from each participating broadband provider was selected to ensure that the data collected would support statistically valid inferences based on a first order analysis of gathered data. Other methodological factors and considerations influenced the selection of the sample size and makeup:

- The panel of EU broadband subscribers was drawn from a pool of over 110,000 volunteers following an on going recruitment campaign that ran from February 2011.
- The volunteer sample was organized with a goal of covering major ISPS in the 27 Member States + Croatia, Iceland and Norway across the available broadband technologies: xDSL, Cable, FTTx.
- For the purposes of creating the sample plan service tiers were split into 4 service tier ‘bands’, as follows: <2 Megabits per second (Mbps), 2<8 Mbps, 8<30Mbps and >=30 Mbps with each speed tier forming an individual sample ‘cell’ against which a target number of volunteers would be selected.
- A target plan for allocation of Whiteboxes was developed based on the market share of participating ISPs. Initial market share information was based principally on data from the European Commission and some commercial sources. However it is worth noting that the distribution of Whiteboxes does not need to replicate the market share in each market as outlined in B.3.3.
- An initial set of prospective participants was selected from volunteers who had responded directly to SamKnows as a result of media solicitations. This pan-European recruitment drive ran from September 2011 to February 2012 and was supported by the European Commission and some commercial sources.


- The sample plan was reviewed and approved by statistical experts from the European Commission, who had the final say in any recruitment methodology and sample plan.
- **It should be noted that unlike other SamKnows studies, ISPs did not participate in recruitment aside from some limited exceptions in the following countries: UK, Holland and Belgium. It is hoped that ISPs will be given the opportunity to participate in future reports.**

B.1.1.4.3 **Panellist Recruitment Methodology**

The following 7-step process is used to recruit a representative panel of volunteer European broadband consumers:

1. Each consumer-volunteer is directed to a web form that is built specifically for each project. Once there they are asked to complete a short form which gives SamKnows a minimum amount of personal information so that SamKnows can determine whether each volunteer fits the sample plan requirements.
2. Once selected the volunteer is then sent an email which asks for further information and requests that the volunteer complete a speed test. This speed test has been developed in-house by SamKnows. On the basis of the information provided by the customer and the results of the speed test, it is then decided whether the consumer is eligible for the next stage.
3. If successful, the volunteer is sent an End User License Agreement that details key considerations such as the responsibilities of the volunteer and SamKnows respectively, data ownership and duration of the project. On completion of the EULA SamKnows will then organize for the dispatch of a Whitebox.
4. A spreadsheet is sent via encryption to our distribution partner of over 4 years. Ours is a long-standing relationship built on shared foundations of trust, expertise and dedication. SamKnows has built a number of automatic systems that the distribution partner can pull data from and also push data to. It is in this way that SamKnows is able to automatically check the progress of Whiteboxes as they are dispatched from the warehouse and ensure they are delivered to the volunteer in a timely fashion.
5. Once they receive confirmation that the Whitebox has been delivered, the SamKnows support team will contact the volunteer in order to provide assistance with installing the Whitebox. Over four years SamKnows has developed a sophisticated support team infrastructure that ensures our volunteers receive the best possible support during installation and throughout the project.
6. Once installed the Whitebox calls 'home' to request the SamKnows test suite. Once received it will start to run tests according to the pre-defined testing schedule. The volunteer is then sent an email automatically by the reporting system with details of how to access the results data from the SamKnows reporting system.
7. The entire process and system has been developed over the course four years and is constantly being refined. As a consequence of this focus and development the volunteer recruitment through to on-boarding and on-going support is an extremely efficient process. This methodology has

been used successfully in the following continents: Asia, North America, South America and Europe.



B.1.4.4 Validation of Volunteers' Service Tier

The methodology employed in this study included verifying each panelist's service tier and ISP against the recorded base (the advertised speeds) of participating ISPs. Initial throughput tests were used to confirm reported speeds.

The broadband service tier reported by each panelist was authenticated in the following way:

- At the time of recruitment, each panelist was required to complete a speed test using the SamKnows web-based speed test. This test provided a rough approximation of the panelist's service tier, which served to identify panelists with targeted demographics, and highlighted anomalies in the panelist's survey response to measured speed.
- At the time the panelist installed the Whitebox, the device automatically ran an IP test to check that the ISP identified by the volunteer was correct.
- The Whitebox also ran an initial test which flooded each panelist's connection in order to accurately detect the throughput speed when their deployed Whitebox connected to a test node.

SamKnows manually completed the following four steps for each panelist:

1. Verified that the IP address was in a valid range for those served by the ISP in question.
2. Reviewed data for each panelist and removed data where speed changes such as tier upgrade or downgrade appeared to have occurred, either due to a service change on the part of the consumer or a network change on the part of the ISP.
3. Identified panelists whose throughput appeared inconsistent with the provisioned service tier.
4. Verified that the resulting downstream-upstream test results corresponded to the ISP-provided speed tiers, and updated accordingly if required.

B.1.4.5 Protection of Volunteers' Privacy

A major concern during this trial was to ensure that panelists' privacy was protected. The panel was comprised entirely of volunteers who knowingly and explicitly opted-in to the testing program. Full opt-in documentation was preserved in confidence for audit purposes.

All personal data was processed in conformity with relevant European laws and in accordance with policies developed to govern the conduct of the parties handling the data. The data were processed solely for the purposes of this study and are presented here and in all online data sets with all personally identifiable information (PII) removed.

B.1.5 **Data Analysis Methodology**

B.1.5.1 **Data Integrity**

As the Whiteboxes ran tests consistently from homes across the European Union, it was important to check the data to ensure that any anomalies were removed. To ensure the integrity of the large amount of data collected, the following protocols were developed:

1. Change of ISP intra-month: found units that changed ISP intra-month (determined by performing daily WHOIS query using the panelist's IP address), and removed data for the ISP on which they spent less time over the course of that month.
2. Change of service tier intra-month: found units that changed service tier intra-month by isolating the difference between the average sustained throughput observed for the first three days in the reporting period from the average sustained throughput observed for the final three days in the reporting period. If a unit was not online at the start or end of that period, then the first/final three days that they were actually online were taken. If this difference was over 50%, the downstream and upstream charts for this unit were individually reviewed. Where an obvious step change was observed (e.g., from 1 Mbps to 3 Mbps), the data for the shorter period was flagged for removal.
3. Removal of any failed or irrelevant tests: removed any failed or irrelevant tests by removing measurements against a pre-defined criteria, for example, measurements against any test node (server) that exhibited greater than or equal to 10% failures in a specific one hour.
4. Removal of any problem units: removed measurements for any unit that exhibited greater than or equal to 10% failures in a particular one hour period (the purpose was to remove periods where units were unable to reach the internet).

B.1.5.2 **Collation of Results and Outlier Control**

All measurement data were collated and stored for analysis purposes as monthly trimmed averages during three time intervals (24 hours, 7:00 pm to 11:00 pm local time Monday through Friday, 12:00 am to 12:00 am local time Saturday and Sunday). Only participants who provided a minimum of one week (seven days) of valid measurements and had valid data in each of the three time intervals were included in the test results. In addition, we dropped the top and bottom 1% of measurements to control for outliers that may have been anomalous or otherwise misrepresentative of actual broadband performance. All statistics were computed on the trimmed data with a minimum sample of 50 reporting Whiteboxes.

The resulting final sample of data for March 2012 was 9,104 participants.

B.1.5.3 Peak Hours Adjusted to Local Time

Peak hours were defined as weekdays between 7:00 pm to 11:00 pm (inclusive) for the purposes of the study. All times were adjusted to the panelist's local time zone. Due to some tests that only took place once every two hours on an individual Whitebox, the period used for aggregating peak performance had to be a multiple of two.

B.1.5.4 Congestion in the Home Not Measured

Download, upload, latency, and packet loss measurements were taken between the panelist's home gateway and the dedicated test nodes. Web browsing measurements were taken between the panelist's home gateway and three popular EU hosted websites. Any congestion within the user's home network is therefore not measured by this study. The web browsing measurements are subject to possible congestion at the content provider's side, although the choice of three highly trafficked websites configured to serve high traffic loads may have mitigated the effects of temporary congestion.

B.1.5.5 Latencies Attributable to Propagation Delay

The speeds at which signals can traverse networks are limited at a fundamental level by the speed of light. While the speed of light is not believed to be a significant limitation in the context of the other technical factors addressed by the testing methodology, a delay of 5 ms per 1000 km of distance travelled can be attributed solely to the speed of light. The geographic distribution and the testing methodology's selection of the nearest test servers are believed to minimize any significant effect. However, propagation delay is not explicitly accounted for in the results.

B.1.5.6 Limiting Factors

A total of 3,065,341,850 measurements were taken across 75,978,173 unique tests.

All scheduled tests were run, aside from when monitoring units detected concurrent use of bandwidth.

Definitions

Technology Splits

Results in sections C, and D are often split by access technology. This report defines these technologies as ‘xDSL’, ‘Cable’ and ‘FTTx’. The services that these encompass are defined as follows:

Technology	Description
xDSL	All residential ADSL, ADSL2+ and SDSL services.
Cable	Residential services delivered by coaxial cable to a cable modem in the user’s premises.
FTTx	Residential fibre-to-the-home and fibre-to-the-cabinet services (including those that use VDSL for the last leg to the home).

Scenario Matrix

Scenario	Metric	Impact
Low Download Speed	Web Browsing	At download speeds under 10Mbps, web browsing increases speed at a linear rate, and then levels out. So if download speed is under 10Mbps, users will notice a drop in web browsing performance
Low Upload Speed	Download Speed	A really low upload speed could negatively impact download speed, as TCP ACKs cannot reach the server fast enough, effectively choking the download.
	Web Browsing	A really low upload speed (~128k) will hamper web browsing.
High but stable latency (>100ms)	Download Speed	Throughput may be affected on very fast lines (100Mbps+) as bandwidth-delay-product becomes a dominating factor.
	Upload Speed	Throughput may be affected on very fast lines (100Mbps+) as bandwidth-delay-product becomes a dominating factor.
	Packet Loss	Loss will likely be higher as there is more time and locations for packets to be lost.
	Web Browsing	Web browsing performance will suffer very noticeably as round-trips are limited by the latency achieved
Very variable latency (high rtt_stddev in curr_udplacency)	Download Speed	Will likely be very variable.
	Upload Speed	Will likely be very variable.
	Packet Loss	Highly variable latency usually accompanies significant packet loss, so expect larger numbers here.
	Jitter	Jitter would likely be very high.
	Web Browsing	Web browsing performance will suffer very noticeably and we may even see some failures due to any associated packet loss.
High packet loss (>5%)	Download Speed	Likely to see highly variable speeds at best, and more realistically we will see lots of speed tests failing (failures>0)
	Upload Speed	Likely to see highly variable speeds at best, and more realistically we will see lots of speed tests failing (failures>0)
	Latency	Latency needn't necessarily be affected, but it

Very high/unstable jitter (>100ms)		probably will be.
	Jitter	Jitter needn't necessarily be affected, but it probably will be.
	Web Browsing	Expect lots of web browsing tests to fail completely, or at least show very poor results.
	Download Speed	Will likely be very variable.
	Upload Speed	Will likely be very variable.
	Latency	Latency will likely be highly variable
	Packet Loss	May or may not be affected.
	Web Browsing	Web browsing performance will suffer very noticeably as round-trips are limited by the latency achieved

Data Dictionary

curr_dns.csv

unit_id	Unique identifier for an individual unit
mtime	Time test finished in UTC
nameserver	Nameserver used to handle the DNS request
lookup_host	Hostname to be resolved
response_ip	Field unused at present
rtt	DNS resolution time in microseconds
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

curr_httpgetmt.csv

unit_id	Unique identifier for an individual unit
mtime	Time test finished in UTC
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for in microseconds
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval
bytes_sec_interval	Throughput at this specific interval (e.g. Throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size (Units: microseconds)
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase.
sequence	The interval that this row refers to (e.g. in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)

curr_httppostmt.csv

unit_id	Unique identifier for an individual unit
mtime	Time test finished in UTC
target	Target hostname or IP address
address	The IP address of the server (resolved by the client's DNS)
fetch_time	Time the test ran for in microseconds
bytes_total	Total bytes downloaded across all connections
bytes_sec	Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of

	the test to the current interval
bytes_sec_interval	Throughput at this specific interval (e.g. Throughput between 25-30 seconds)
warmup_time	Time consumed for all the TCP streams to arrive at optimal window size (Units: microseconds)
warmup_bytes	Bytes transferred for all the TCP streams during the warm-up phase.
sequence	The interval that this row refers to (e.g. in the US, sequence=0 implies result is for 0-5 seconds of the test)
threads	The number of concurrent TCP connections used in the test
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)
curr_udpjitter.csv	
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
packet_size	Size of each UDP Datagram (Units: Bytes)
stream_rate	Rate at which the UDP stream is generated (Units: bits/sec)
duration	Total duration of test (Units: microseconds)
packets_up_sent	Number of packets sent in Upstream (measured by client)
packets_down_sent	Number of packets sent in Downstream (measured by server)
packets_up_rcv	Number of packets received in Upstream (measured by server)
packets_down_rcv	Number of packets received in Downstream (measured by client)
jitter_up	Upstream Jitter measured (Units: microseconds)
jitter_down	Downstream Jitter measured (Units: microseconds)
latency	99th percentile of round trip times for all packets
successes	Number of successes (always 1 or 0 for this test)
failures	Number of failures (always 1 or 0 for this test)
location_id	Please ignore (this is an internal key mapping to unit profile data)
curr_udplatency.csv	
	UDP based
unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	Target hostname or IP address
rtt_avg	Average RTT in microseconds
rtt_min	Minimum RTT in microseconds
rtt_max	Maximum RTT in microseconds
rtt_std	Standard Deviation in Measured RTT in microseconds
successes	Number of successes (note: use failures/(successes+failures)) for packet loss

failiures	Number of failures (packets lost)
location_id	Please ignore (this is an internal key mapping to unit profile data)

curr_webget.csv

unit_id	Unique identifier for an individual unit
dtime	Time test finished in UTC
target	URL to fetch
address	IP address connected to to fetch content from initial URL
fetch_time	Sum of time consumed to download Html content and then concurrently download all resources (Units: micorseconds)
bytes_total	Sum of HTML content size and all resources size (Units : Bytes)
bytes_sec	Average speed of downloading HTML content and then concurrently downloading all resources (Units: bytes/sec)
objects	Number of Resources (images, css etc) downloaded
threads	Maximum number of concurrent threads allowed
requests	Total number of HTTP requests made
connections	Total number of TCP connections established
reused_connections	Number of TCP connections re-used
lookups	Number of DNS lookups performed
request_total_time	Total duration of all requests summed together, if made sequentially
request_min_time	Shortest request duration
request_avg_time	Average request duration
request_max_time	Longest request duration
tffb_total_time	Total duration of the time-to-first-byte summed together, if made sequentially
tffb_min_time	Shortest time-to-first-byte duration
tffb_avg_time	Average time-to-first-byte duration
tffb_max_time	Longest time-to-first-byte duration
lookup_total_time	Total duration of all DNS lookups summed together, if made sequentially
lookup_min_time	Shortest DNS lookup duration
lookup_avg_time	Average DNS lookup duration
lookup_max_time	Longest DNS lookup duration
successes	Number of successes
failures	Number of failures
location_id	Please ignore (this is an internal key mapping to unit profile data)

B.3

Further information regarding sampling methodology

B.3.1

Identifying the Test Variables

We begin by defining the single factor to be measured and analysed; in statistical sampling theory this is known as the “dependent variable”. The dependent variable can be best described as an output that varies according to the size and type of certain inputs, where the aim is to measure accurately the impact on the dependent variable of changes in the input variables. For example, if we are trying to reach conclusions about download speeds (download speed is our dependent variable) then we may consider the impact of factors such as geography and ISP.

Having established the dependent variable we are measuring, we then look for the other factors that might influence the dependent variable. These are called “explanatory variables” because they explain changes in the dependent variable we are seeking to measure. In some cases we will also seek to distinguish between “primary” and “secondary” explanatory variables: primary explanatory variables are the inputs that we explicitly intend to test; secondary variables will not be tested but may also have an influence on the dependent variable.

These will determine how the sample is constructed. We use the primary explanatory variables to define the subgroups that the sample will be divided into - these subdivisions are known as ‘strata’.

These ‘primary’ factors are then used to break down the sample into subpopulations. These subdivisions are mutually exclusive: such as a binary split into two different broadband speed bands (above or below 2Mb, say).

In the European Commission study the test variables were as follows: country and technology. To measure these we therefore recruited a voluntary panel of European broadband customers across Europe which was consistent with this criteria. Please note that over the initial 6 months of the project, the panel composition changed to accommodate a revision to the test variables.

B.3.2

Subdivisions of the Population

Next we define quotas for each sub-group, so that the number and characteristics of participants in each subdivision are known in advance. “Quota sampling” allows us to allocate participants in these subgroups in a proportion that is highly representative of the wider population.

This means that, in effect, we have theoretically broken down our population into subdivisions (for example, based around the key factors that we believe explain download speed) and we will now reconstruct a sample of this population by recruiting participants directly into particular subdivisions dependent on their characteristics. For instance, for a subgroup which has a speed band below 2Mb

and is in a specific area (member state), we will screen and recruit volunteers into this subdivision by using only the participants who exactly match this sample criteria.

Thus, unlike random sampling, quota sampling ensures that all subdivisions of the population are represented in the sample. We have the additional benefit of being able to use a smaller sample and yet also potentially gaining much greater accuracy than a purely random sample may achieve because our sample is - by its nature - constructed to be representative of the whole population. Conclusions from relatively small samples can then be, very reliably, extrapolated to the larger population of internet users from which the sample was drawn.

B.3.3 Sample Parameters

In determining the size of the sample, we must first consider that our desired confidence. For example, a confidence level of 95% means that if we were to repeat our test many times, we would expect the 'true' value of our dependent variable to fall within the interval we actually observe 95% of the time. This is a very standard and accepted level of confidence used in order to make statistically significant conclusions.

The size of this interval is, in turn, determined by the margin of error. If we are measuring the impact of location on download speed for instance, then we would measure the download speed subject to an error margin. There is no universally accepted error margin, as the appropriate level will depend heavily on the nature and distribution of the dependent variable, which is likely to vary for different dependent variables. For example, testing download speed and testing latency may require completely different margins of error. Furthermore, if a variable has a wide dispersion (it is spread out over a larger range) then we may be more inclined to accept a larger error margin. We must also consider that choosing the error margin is a trade off between data accuracy, and the necessary sample size. In general, for a given confidence level, the smaller the margin of error is then the larger the sample. A minimum number of participants for each subgroup may also be necessary (since each subgroup in the sample must be representative of the corresponding population subgroup), in order to create accurate data and to conduct analysis that is statistically significant.

To ensure that this panel is consistent with others we have previously built to complete other studies we adopted identical sample parameters. Initially we used the model of the Federal Communications Commission (USA) which dictates that a sample must exceed 25 Whiteboxes (panellists) to be valid. However, after consultation it was decided to increase this minimum level to 50 to be consistent with our studies on behalf of Ofcom (UK). This number ensures a confidence interval of 98%.

B.3.4 Participant Recruitment

Volunteers are then recruited to meet this previously defined quota and sample size. Many more participants are recruited than are actually desired in the end

sample. Our previous sampling experience dictated that typically ten initial participants should be recruited for each one in the target final sample size. This is necessary because of high rates of volunteer attrition, common to many types of sampling, in addition to the exclusion of unsuitable participants (such as those who are prohibitively long distances from termination points).

B.3.5 Statistical Analysis Considerations

Recall that we know the dependent variable is likely to be determined by a wide range of explanatory variables, not just the primary factors such as ISP or location that are explicitly being tested. It is critical to also acknowledge the secondary explanatory variables. These are all the other factors that are not being tested but that we still expect to have a significant impact on the dependent variable.

Whilst these secondary factors may not be of direct interest, their impact must still be taken into account in order to avoid bias (known as ‘omitted variable’ bias in statistical analysis) in the data, and therefore draw valid conclusions. Crucially, this does not necessarily entail further sub-divisions of the sample; simply using proportional weighting in the data analysis phase can control for its impact.

For example, in order to isolate the effects of location (a primary explanatory variable) on download speed (our dependent variable), we may identify different technologies, and the distance to a termination point, as secondary explanatory variables. So whilst we may not explicitly test for the effect of technology, we may still reasonably expect it to have a considerable impact on download speeds. This impact can be ‘controlled’ for by giving more weight to more popular technologies; the weighting is done in direct proportion to the (known) ratio of technologies used by the population. In this way, the sample data remains perfectly representative of the overall population. Then an accurate relationship between location and download speed can be established whilst effectively holding the influence of the type of technology constant.

The length of copper wire between a broadband user and the server can introduce distortions to xDSL performance, known as the “last-mile” effect. To provide for like-for-like comparison in the Ofcom study, data was weighted for the xDSL operators to normalise for distance from the exchange. Weighting by line length specifically addressed differences in performance that could have been introduced as a consequence of line length.

This adjustment was important in order to ensure that an ISP with nationwide coverage is not represented as having poor performance compared to an ISP focused on more densely populated areas simply because it has customers with typically longer line lengths between the premises and exchange.

Please note that for the first European Commission study, the data required to weight the sample according to the distance from the termination point was unavailable and therefore it was not possible to weight the data based on this dependent variable. As stated in this report, it is our expectations that future

reports will include weighting by distance (line length) subject to the participation of the local ISPs.

However, the European Commission study (2012) found that the vast majority of variation in overall performance across member states was attributable to differences in the country's technology composition. It is still preferable to include distance from the exchange as a secondary explanatory variable, wherever data allows. However, inclusion can reasonably be seen as for the purpose of providing an extra layer of control and accuracy rather than explaining performance.

The xDSL metrics that may be affected by distance are primarily download and upload speed, latency and packet loss. The greater the distance to the server, the higher (i.e. worse) the effective performance for these measures is likely to be. Not accounting for distance risks introducing an upward bias for these metrics. To illustrate, Malta, Iceland and Cyprus and Spain were all found to have noticeably higher than average latency. Although one may initially suspect this to be related to an upward bias, looking at the results in comparison with the number and location of servers, it is apparent that Spain has high latency despite multiple measurement servers located within the country. With many servers, distance between the user and server will generally be quite short; theoretically a bias is not expected. The same is true for Italy with respect to packet loss, which has several servers in major cities and yet exhibits very a high packet loss of 1.76% at peak times. It would therefore be ineffective, or even counterproductive, to try to reverse this bias in the data by scrubbing anomalies with high latency or packet loss, as not all such performances appear to be due to bias from distance.

Further, Spain's latency for cable technology is above average to a similar degree as Malta and Iceland, suggesting that even if there is some upward bias, it is generally likely to be within the range of regular, nonbiased observations. Thus, with little or no explanatory power, there is not a risk of omitted variable bias and the study's results will still be valid.

Finally, in the short term, not weighting by distance may introduce some slight inaccuracy for xDSL service providers. In the long run, as European users continue the shift away from xDSL towards superior FTTx and cable technologies, the potential impact of this bias will fall.

B.3.6 Data Validation

The process of controlling for these secondary explanatory variables forms a part of our "data validation" phase, which ensures that the data collected from the sample that we report on is both accurate and unbiased. Filtering the raw data also forms a large part of validation, whereby we control for anomalies or outliers which might otherwise skew the results. Following a common precaution for controlling statistical outliers, we typically exclude the top and bottom 1% of the data from our analysis. This exclusion rate to remove outliers is chosen based on SamKnows' past and current experience and is sufficient to control for outliers. It is unnecessary for the exclusion rate to be any higher than 1%.

This method serves to remove the most extreme data points that might otherwise skew the results and thereby misrepresent the broadband performance for the typical user. Note that, by this stage of data validation, steps have already been taken to collect data expected to be accurate and highly representative. Trimming the collected data by the top and bottom percentile therefore acts as an additional final safety net.

An early example of control is in constructing sampling plans. Within country and technology, speed tiers used by less than 5% of the country's population are not included. In this way, only material speed tiers are tested, as well as avoiding potential bias from cells with very small numbers of participants. For example, three speed tiers are excluded for this reason for Belgium, but the sample still covers 96.7% of the country.

Other techniques at the validation stage help control for anomalous events such as participants switching ISP, or broadband tier within an ISP. Further, only participants providing at least one week of valid measurements, and who had valid data in all three daily time intervals were accepted in the confirmed data set. Moreover, failures are excluded from the raw data (unless of course, we are measuring the dependent variable as the number of failures).

B.3.7 Impact of Unaccounted Variables (IPTV)

The SamKnows measurement solution is designed to control for all material variables. Although in some cases IPTV traffic does not directly pass through the Whitebox, the current version of the SamKnows system is able to infer and control for the impact. The key to understanding the impact of IPTV is an ability to profile the performance of an IPTV-enabled internet connection. It is then possible to spot for performance variation that is as a consequence of IPTV, rather than network congestion. This is something that is being developed by SamKnows analysts, with the intention of this functionality being built-in to the user reporting. Note that the impact of IPTV on broadband performance varies according to provider and package since not all IPTV services share bandwidth with normal internet traffic.

C EU Level Analysis¹

C.1 Key Performance Indicators

C.1.1 Download Speed

Figure EU.1-1 and Figure EU.1-2 depict actual download speed as a percentage of advertised speed over peak periods and 24-hour periods, split by access technology. Cable and FTTx delivered 91% and 84% of advertised download speed respectively. xDSL achieved a significantly lower figure, delivering 63% of advertised speeds. All technologies experienced small declines in their performance during peak periods.

Please note that the figures below are not derived by dividing the average actual speed through the average advertised speed. They are instead computed on a per-panelist basis and averaged to form the overall figure. This approach is effectively a mean of ratios rather than a ratio of means.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Actual Speed	7.19	7.43	33.08	34.13	41.02	42.46	19.47	20.12
Advertised Speed	12.87	12.88	37.50	37.50	50.58	50.65	25.60	25.61
Actual/Advertised Speed	63.3%	65.4%	91.4%	94.6%	84.4%	87.2%	74.0%	76.0%

Figure EU.1-1: Actual Peak and 24-hour Period Download Speed as a Percentage of Advertised Speed, by technology (higher is better)

¹ EU refers to the average of all the countries included in the sample, i.e. EU27 countries & Croatia, Iceland and Norway




Figure EU.1-2: Actual Peak and 24-hour Period Download Speed as a Percentage of Advertised Speed, split by technology (higher is better)

Figure EU.1-3 below shows the actual download speed as a percentage of advertised download speed split by time of day. All technologies exhibit the same pattern, where download speed performs better during the morning hours before declining slightly during the afternoon and evening. It is during the evening hours that consumers are more likely to make use of their broadband services.




Figure EU.1-3: Actual Download Speed as a Percentage of Advertised Speed, split by hour of day and technology (higher is better)

Figure EU.1-4 below plots actual download speed split by time of day and access technology. As was the case in Figure EU.1-2, all technologies exhibit the same behavioural pattern where speed experiences a slight decline in the afternoon and

evening hours. This decline is most evident for the cable and FTTx technologies, where access speeds are often far higher than that of xDSL. It is worth noting that whilst cable services delivered a higher percentage of advertised speed (see Figure EU.1-2), FTTx outperforms it in absolute terms.




Figure EU.1-4: Actual Download Speed split by hour of day and technology (higher is better)

C.1.2 Actual Download Speed split by hour of day and technology

Figure EU.1-5 presents a cumulative distribution plot of download speed as a percentage of advertised speed. This chart shows the percentage of consumers who receive at least a certain level of their advertised broadband speed. In other words, the chart gives an indication of what proportion of consumers are receiving the advertised speeds. Earlier charts in this report have focused on averages alone, but these can mask high levels of inconsistency; this helps show if there is significant spread amongst the sample.

For example, one technology may deliver 80% of advertised to all of its users at all times, and another technology may deliver anywhere between 60% and 100% of advertised. Both may produce averages of 80% and thus be indistinguishable in charts earlier in the document. A cumulative distribution plot would highlight these differences clearly.

We can see from Figure EU.1-5 that 80% of cable consumers receive at least 84% of advertised speed or better, FTTx consumers receive 71% and xDSL consumers receive only 35%. The chart overall shows that most cable and FTTx consumers receive above 80% of the advertised rates, whereas there is a much wider distribution among xDSL consumers, with only half receiving download speeds above 80% of advertised speeds. It is worth noting that xDSL performance decreases with the length of the copper line connecting the consumer to the termination point (usually a telephone exchange or cabinet) as outlined in B.3.5




Figure EU.1-5: Cumulative Distribution of Download Speed as a Percentage of Advertised Speed, split by technology

c.1.3 Upload Speed

Figure EU.1-6 and Figure EU.1-7 show the actual upload speed as a percentage of advertised speed over peak and 24-hour periods across the types of technology considered for this study. All technologies exhibited actual upload speeds over 80% of their advertised speeds, with cable delivering just over 100% of advertised speed during the 24-hour period and just below it during peak periods. As with download speeds, all technologies experienced a small decline in their performance during peak periods.

As is the case with download speed, the figures below are not derived by dividing the average actual speed through the average advertised speed. They are instead computed on a per-panelist basis and averaged to form the overall figure.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Actual Speed	0.69	0.69	3.68	3.70	19.8	20.92	5.02	5.19
Advertised Speed	0.88	0.88	3.80	3.79	25.28	25.37	6.20	6.15
Actual/Advertised Speed	82.8%	83.1%	99.4%	100.2%	88.1%	90.5%	88.0%	88.0%

Figure EU.1-6: Actual Peak and 24-hour Period Upload Speed results as a Percentage of Advertised Speed, by technology (higher is better)

Figure EU.1-7 and EU.1-8 show the actual upload speed as a percentage of advertised speed split by technology and by time of day. Cable and xDSL based services deliver stable upstream throughput throughout the day, whilst FTTx exhibits a small but steady decline into the evening hours. All technologies deliver marginally lower speeds during peak hours.




Figure EU.1-7: Actual Peak and 24-hour Period Upload Speed as a Percentage of Advertised Speed, split by technology (higher is better)




Figure EU.1-8: Actual Upload Speed as a Percentage of Advertised Speed split by hour of day and technology (higher is better)

Figure EU.1-9 shows actual upload speed split by type of technology and by hour of day. The same behaviour that was observed in Figure EU.1-8 is mimicked here; cable and xDSL remain stable throughout the day, whilst FTTx sees a small decline. Of course, this chart clearly demonstrates that FTTx delivers a far higher average upstream throughput than the other technologies. So even with a small dip during the evening hours, it far outperforms the alternatives.




Figure EU.1-9: Actual Upload Speed split by hour of day and technology (higher is better)

Figure EU.1-10 is a cumulative distribution plot of upload speed as a percentage of advertised speed. As was the case with download speed, there is a very even distribution among cable consumers, where 80% of cable users receive at least 95% of advertised speed or more. FTTx and xDSL consumers, however, are spread over a wider distribution, although it is to a far lesser extent for xDSL consumers than was seen for download speed, where around 60% of xDSL consumers receive at least 80% of advertised speed. This tighter spread will be due to the asymmetric nature of broadband services; upstreams are typically provisioned at a far lower rate than the downstreams.




Figure EU.1-10: Cumulative Distribution of Upload Speed as a Percentage of Advertised Speed, split by technology

C.1.4 Latency

Figure EU.1-11 and Figure EU.1-12 display actual average latency over the peak and 24-hour periods across the various types of broadband technology. Average latency proved to be highest for xDSL users at nearly 40ms during peak and 24-hour periods. All forms of access technology experience a rise in round-trip

latency during peak periods, although FTTx experiences the smallest change.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Average Latency (ms)	39.94	38.46	24.87	23.32	22.01	21.58	33.11	31.82

Figure EU.1-11: Peak period and 24-hour Average Latency results, by technology (lower is better)




Figure EU.1-12: Peak period and 24-hour Average Latency, split by technology (lower is better)

Figure EU.1-13 shows average latency split by hour of day and technology. As indicated in Figure EU.1-12, cable and xDSL technologies experience a more elevated increase in latency during peak hours, but this is still very modest and would be unnoticeable for almost all real world use cases.




Figure EU.1-13: Average Latency, split by hour of day and technology (lower is better)

Figure EU.1-14 is a cumulative distribution plot of average latency. This differs from previous CDF plots in that it does not record the percentage of advertised speed. This is because ISPs typically do not advertise latency for their broadband products (except in very special cases), as it can vary so widely depending on the host the user is communicating with. For the purposes of this chart, and all subsequent charts, the actual value will be plotted instead.

Figure EU.1-14 shows that both cable and FTTx technologies have similar distributions of average latency among its customers, where 60% of both technologies are delivering latencies of 25ms or less. This is to be expected, as both of these access technologies are less affected by the length of the ‘last-mile’ cable than xDSL. 60% xDSL consumers are receiving 45ms or less.




Figure EU.1-14: Cumulative Distribution of Average Latency, split by technology (lower is better)

C.1.5 **Packet Loss**

Figure EU.1-15 and Figure EU.1-16 show average packet loss over the peak and 24-hour periods across different access technologies. Of the three, xDSL displayed the largest increase in packet loss during peak hours. Packet loss during the 24-hour period also proved to be higher for xDSL users compared to cable or FTTx services. Cable and FTTx technologies experienced very small increases in packet loss during peak hours.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Packet Loss (%)	0.7%	0.4%	0.3%	0.3%	0.3%	0.3%	0.5%	0.4%

Figure EU.1-15: Peak period and 24-hour packet loss, split by technology (lower is better)




Figure EU.1-16: Peak period and 24-hour packet loss, split by technology (lower is better)

Figure EU.1-17 shows average packet loss split by hour of day and access technology. xDSL delivers a comparable level of packet loss to cable and FTTx services during the hours of midnight to 10am, but then starts to rise during the daylight and evening hours. Cable and FTTx technologies remained relatively stable throughout, experiencing only slight increases during peak hours.




Figure EU.1-17: Packet Loss, split by hour of day and technology (lower is better)

Figure EU.1-18 presents a cumulative distribution plot for packet loss. Very few consumers experience high levels of packet loss, with 90% of xDSL users exhibiting packet loss of 2% or less. 90% of cable and FTTx users experience at most 1% packet loss. The chart below demonstrates that the majority of consumers from all given technologies are showing very low packet loss figures.




Figure EU.1-18: Cumulative Distribution of Packet Loss, split by technology

C.1.6 DNS Resolution

Figure EU.1-19 and Figure EU.1-20 shows DNS resolution time throughout a 24-hour period and during peak hours. DNS resolution time tends to be much higher for xDSL than for cable and FTTx as it is directly affected by the round-trip latency of the underlying technology. Cable and FTTx technologies demonstrated a lower amount of time taken to resolve DNS queries, with FTTx almost completely unaffected by peak hour traffic, as was the case with latency.

This close correlation with the round-trip latency figures suggest the ISPs' DNS servers are – in general – operating effectively and not adding significant delay to end users' DNS queries.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
DNS Resolution (ms)	40.09	39.37	22.79	21.57	20.53	20.48	32.44	31.73

Figure EU.1-19: Peak Period and 24-hour DNS Resolution Time, split by technology (lower is better)




Figure EU.1-20: Peak Period and 24-hour DNS Resolution Time, split by technology (lower is better)

Figure EU.1-21 below shows DNS resolution time split by hour of day and technology. As was seen in Figure EU.1-20, FTTx is almost completely unaffected by peak hours, maintaining a very stable pattern throughout the day. Cable and xDSL, on the other hand, show small increases during peak periods. On the whole, DNS resolution results for all technologies closely match their latency results.




Figure EU.1-21: DNS Resolution Time, split by hour of day and technology (lower is better)

The cumulative distribution plot in Figure EU.1-22 below shows that just over half of cable and FTTx users experience resolution times of 18ms or less. On the whole, DNS resolution time is lower for FTTx users than it is for cable users. The chart shows that 80% of xDSL users experience a DNS resolution time of 53ms or better.




Figure EU.1-22: Cumulative Distribution of DNS Resolution Time, split by technology

DNS Resolution Failure Rate

Figure EU.1-23 and Figure EU.1-24 demonstrate peak period and 24-hour DNS resolution failure rate across the three different common technologies. No individual technology displayed any significant differences between each other. All technologies experienced a slightly higher failure rate during peak periods. Once again, FTTx displayed the lowest difference between the 24-hour figures and peak period performance.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
DNS Failure Rate (%)	0.5%	0.4%	0.5%	0.4%	0.4%	0.4%	0.4%	0.4%

Figure EU.1-23: Peak period and 24-hour DNS Resolution Failure Rate, split by technology (lower is better)




Figure EU.1-24: Peak period and 24-hour DNS Resolution Failure Rate, split by technology (lower is better)

Figure EU.1-25 shows DNS Resolution Failure Rate split by hour of day and technology. Interestingly, failure rate was shown to be much higher in the morning than during the evening hours, although figures still exhibited a certain increase during peak hours, but not to the same extent as was seen during the early hours. This pattern was observed across all technologies, and suggests that an upstream nameserver (most likely one of the target domain names) was faulty for a short period. As seen in Figure EU.1-24, failure rate was lowest for FTTx users as well as the increase during peak hours.




Figure EU.1-25: DNS Resolution Failure Rate, split by hour of day and technology (lower is better)

Figure EU.1-26 displays the cumulative distribution plot of DNS resolution failure rate for each of the varying technologies in Europe. The chart shows that, similar to packet loss, high DNS resolution failure rates were very rare for the majority of users for all given types of technologies. 70% of users experienced less than 1% failure rate and better.




Figure EU.1-26: Cumulative Distribution of DNS Resolution Failure Rate, split by technology

C.1.7 Web Browsing

Figure EU.1-27 and Figure EU.1-28 depict the average web page fetching times across the varying types of technologies during peak periods and 24-hour periods. It is important to note that this test measures against real websites (Google, Facebook and YouTube), which are geographically distributed across Europe. The test measures the network loading time, not the page rendering time, which will vary by browser and computer performance.

xDSL technology proved to have the slowest web browsing speed, nearly half as fast as cable and FTTx. All technologies experienced a modest increase in loading times during peak hours, with xDSL experiencing the largest increase during said period. Loading times proved not much different between cable and FTTx products, with FTTx users experiencing slightly better web browsing speeds. This small difference will almost certainly be driven by the significantly higher throughput found with FTTx, which offers diminishing returns beyond approximately 10Mbps.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Web page loading time (ms)	1.61	1.55	0.86	0.82	0.83	0.80	1.29	1.24

Figure EU.1-27: Peak Period and 24-hour Webpage Loading Time, split by technology (lower is better)




Figure EU.1-28: Peak Period and 24-hour Webpage Loading Time, split by technology (lower is better)

Figure EU.1-29 below shows the average web page loading time by hour of day and technology. All technologies displayed very similar behaviours as they did for latency. Here, peak hour increases are more evident for all technologies, particularly for xDSL products, which experienced the highest increase. Web browsing speed is more stable for cable and FTTx users, with no significant changes observed during peak hours.




Figure EU.1-29: Average Webpage Loading Time, split by hour of day and technology (lower is better)

Figure EU.1-30 is the cumulative distribution plot for webpage loading times split by technology. The figure below shows that 90% of cable and FTTx consumers experienced loading times of 1 second or less. 90% of xDSL consumers faced a wait of up to 3 seconds. The distribution for all technologies is very thin as the majority are experiencing good web browsing speeds.




Figure EU.1-30: Cumulative Distribution of Webpage Loading Time, split by technology

VOIP

Downstream VoIP Jitter

Figure EU.1-31 and Figure EU.1-32 show peak period and 24-hour results for downstream jitter for the varying types of technology. 24-hour downstream jitter proved to be identical for cable and FTTx technologies, although cable experienced a greater increase in jitter during peak hours. This difference, however, was not significant. xDSL products displayed the greatest difference between 24-hour and peak period figures as well as having the highest downstream jitter figures across the three technologies.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Downstream Jitter (ms)	1.03	0.76	0.62	0.46	0.54	0.46	0.84	0.64

Figure EU.1-31: Peak period and 24-hour Downstream VoIP Jitter, split by technology (lower is better)




Figure EU.1-32: Peak period and 24-hour Downstream VoIP Jitter, split by technology (lower is better)

Figure EU.1-33 shows downstream jitter by hour of day and technology. All technologies display very similar patterns throughout the day, with jitter steadily rising before reaching their peak during the evening, when most consumers are online. Increase in downstream jitter is most evident for xDSL products, whereas cable and FTTx proved to be more stable, with FTTx being the least affected by peak hour traffic.




Figure EU.1-33 Downstream VoIP Jitter, split by hour of day and technology (lower is better)

Figure EU.1-34 presents the cumulative distribution of downstream VoIP jitter for the different types of technology in Europe. As has been the case of latency and browsing speeds, there is a tight distribution for all technologies, particularly cable and FTTx, where 80% of consumers experience less than 1ms of downstream jitter or better. 80% of xDSL consumers, also typically see less than 1ms or better.




Figure EU.1-34: Cumulative Distribution of Downstream VoIP Jitter, split by technology

Upstream VoIP Jitter

Figure EU.1-35 and Figure EU.1-36 show upstream jitter for the varying types of technology in use over peak and 24-hour periods. Cable products proved to have the highest upstream jitter compared to both xDSL and FTTx, as well as experiencing the greatest peak period increase. FTTx, by contrast, did not show upstream jitter beyond 1ms and was not greatly affected by peak hour traffic. xDSL also proved to be interesting, with figures being approximately just 1ms greater than jitter from FTTx products.

The reason for cable services exhibiting higher jitter (particularly for the upstream) is due to the fact that they are based upon the concept of TDMA (Time Division Multiple Access). Effectively the modem's time is divided into slots, during which it can either send or receive data. So if the modem is busy whilst the user tries to send a packet then that packet will have to wait in a queue until there is an opportunity to send it. This can result in small but frequent variations in packet delays, which is effectively what jitter represents.

It is important to note that whilst upstream jitter is often noticeably higher for cable networks, its level is often still so low that it would be unnoticeable for almost all use cases. For example, most Voice over IP (VoIP) phones have a de-jitter buffer of at least 25ms, meaning jitter under 25ms would not effect call quality at all.

The above explanation does not account for cable's more noticeable rise in upstream jitter during peak periods, which is most likely caused by increased usage.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	Overall Peak	Overall 24hr
Upstream Jitter (ms)	2.05	1.82	4.40	3.26	0.97	0.84	2.39	1.97

Figure EU.1-35: Peak period and 24-hour Upstream VoIP Jitter, split by technology (lower is better)




Figure EU.1-36: Peak period and 24-hour Upstream VoIP Jitter, split by technology (lower is better)

Figure EU.1-37, which shows upstream jitter by hour of day and technology, makes it very clear that cable products experienced the greatest increase in jitter during peak hours, indicating packets are being received at less regular intervals. xDSL and FTTx showed a much more stable pattern, by contrast, with only a slight change during peak hours, as indicated in Figure EU.1-36.




Figure EU.1-37: Upstream VoIP Jitter, split by hour of day and technology (lower is better)

Figure EU.1-38, the cumulative distribution plot of upstream jitter by each type of technology, shows that distribution is tight for all technologies, although it is cable users who exhibit the widest distribution in this instance. 80% of cable users exhibit approximately 6ms of upstream jitter or better, whereas xDSL and FTTx users display figures already at 1ms or less.




Figure EU.1-38: Cumulative Distribution of Upstream VoIP Jitter, split by technology

C.2

Comparison with the United States²

C.2.1

Download Speed

Figure EU.1-39 below shows the average actual and advertised download speeds in Europe and the USA for each of the technologies studied. Actual download speed in Europe is typically better for all technologies, although it can be observed that actual speeds in the USA are much closer to their advertised speeds, indicating that there are significant differences in the way that broadband products are marketed and sold between the USA and Europe.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual vs Advertised	USA Advertised (Mbps)	USA Actual (Mbps)	USA Actual vs Advertised
xDSL	12.87	7.19	63.3%	6.19	5.29	85.4%
Cable	37.50	33.08	91.4%	17.21	17.03	99.0%
FTTx	50.58	41.02	84.4%	25.92	30.20	116.5%

Figure EU.1-39: Comparison of Average Actual and Advertised Download speed between Europe and the USA, split by technology

C.2.2

Upload Speed

Figure EU.1-40 shows the average actual and advertised upload speeds for the varying types of technology in Europe and the USA. xDSL and FTTx technologies performed better in the USA than in Europe, exceeding actual upload speed in Europe significantly. Upload speed for cable services, however, was superior in Europe in absolute terms. However, in all cases, actual upload speed exceeded advertised speed in the USA, whereas they were noticeably below advertised speeds in Europe.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual vs Advertised	USA Advertised (Mbps)	USA Actual (Mbps)	USA Actual vs Advertised
xDSL	0.88	0.69	82.8%	0.87	0.89	102.6%
Cable	3.80	3.68	99.4%	2.41	2.64	109.8%
FTTx	25.28	19.80	88.1%	22.03	23.18	105.2%

Figure EU.1-40: Comparison of Average Actual and Advertised Upload speed between Europe and the USA, split by technology

² Data taken from Measuring Broadband America - July 2012 - <http://www.fcc.gov/measuring-broadband-america>

C.2.3 Latency

Figure EU.1-41 provides a comparison of latency figures between Europe and the USA, split by technology. With the exception of FTTx, Europe delivered lower average latency than the USA. The difference was not very significant, suggesting that the technology is deployed in a similar manner in both regions. The fact that FTTx produced a lower latency in the USA is to be expected too, as the FTTx sample in the USA was dominated by an FTTP service, whilst Europe's sample included a lot of VDSL users (FTTC).

Technology and Country	Europe	USA
xDSL (ms)	39.94	43.43
Cable (ms)	24.87	26.09
FTTx (ms)	22.01	18.40

Figure EU.1-41: Comparison of Latency between Europe and the USA, split by technology

C.2.4 Packet Loss

Figure EU.1-42 compares packet loss by technology between Europe and the USA. In all cases, the USA exhibited noticeably lower packet loss than Europe. In real world terms, the difference is not significant enough to be noticeable to users.

Technology and Country	Europe	USA
xDSL (%)	0.66%	0.27%
Cable (%)	0.30%	0.13%
FTTx (%)	0.34%	0.13%

Figure EU.1-42: Comparison of Packet Loss between Europe and the USA, split by technology

D Comparison Between Countries

D.1 Key Performance Indicators³

D.2 Download and Upload Speeds

D.2.1 Download

Download throughput is the metric most commonly associated with broadband performance. It is also the one that ISPs typically advertise their products with. As such, it receives a large amount of attention from regulators and ISPs when marketing their products.

In order to provide comparability between countries and technologies, which often have vastly different performance characteristics, results here are presented as a percentage of advertised. This also allows the reader to easily determine how accurate the marketing claims of ISPs in a certain country are.

Figures EU.2-1, EU.2-2 and EU.2-3 below present download speed as a percentage of advertised speed for each country considered in this study for xDSL, Cable and FTTx respectively. Cable services deliver close to the advertised speed in all countries in which they are sold at peak hours. FTTx also fares generally well. In the case of xDSL we see a much wider spread, with the UK and France only achieving just over 40% of advertised – far below the average of 60%. Other large, developed countries such as Germany do not suffer in the same way. This suggests significant differences in the markets; most likely from the advertising practices used.

³ The data has not been weighted, however the data has been trimmed as per section B.1.4.2




Figure EU.2-1: Actual Download Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, split by country




Figure EU.2-2: Actual Download Speed of cable technology as a Percentage of Advertised Speed during peak periods, split by country




Figure EU.2-3: Actual Download Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, split by country

As observed earlier, countries where cable services are prevalent tend to achieve figures close to their advertised rates.

Whilst xDSL is almost ubiquitous, with Latvia, Romania, Belgium and Bulgaria not having large enough representative samples, it is also far more likely to deliver a lower percentage of advertised performance than cable or FTTx. This is due to access speed degrading as with increasing copper loop lengths. Countries such as the UK and France, which feature very close to the bottom of the table, achieved far higher percentages for cable and FTTx access technologies.

Figures EU.2-4, EU.2-5 and EU.2-6 below show the actual speed achieved in each country for each technology. The spread in these figures are noticeably wider, the Nordic and Eastern European states dominating throughput thanks to their deployment of FTTH.

Interestingly, France and the UK deliver much closer to the xDSL average (France slightly beats the average) in this measure. The situation is driven by differences in the way that ISPs advertise broadband products – typically xDSL-based products. In the UK and France, for example, xDSL products are predominantly advertised with a single headline speed (e.g. 20 Mbps). Customers whose copper phone lines mean that they can only receive a fraction of that speed will still be sold that product. Other countries will offer a wider spectrum of products and may adopt policies that they will not sell customers products that they cannot possibly achieve full speed on.




Figure EU.2-4: Actual Download Speed of xDSL technology during peak periods, split by country




Figure EU.2-5: Actual Download Speed of cable technology during peak periods, split by country




Figure EU.2-6: Actual Download Speed of FTTx technology during peak periods, split by country

D.2.2 Upload

The vast majority of European ISPs provide an asymmetric broadband service, with download throughput typically being many multiples of the upload throughput. Historically this has made sense to ISPs as users consumed (downloaded) far more content than they produced and shared (uploaded).

However, upload throughput is becoming an increasingly important metric as more and more users upload photos, videos, and make use of online storage services. Many ISPs nowadays are offering higher speed upload services in recognition of this change.

As with download throughput, the figures for upload are initially presented as a percentage of the advertised rate to provide comparability.

The figures below show upload speed as a percentage of advertised speed by country for each type of technology. The key observation here is that the average percentage of advertised is considerably higher than it was for the download metric. This is most likely due to the asymmetry of the rates (a service needs to handle less traffic in order to deliver a higher percentage in the upstream direction as the rates are lower). For xDSL based services this is critical, as the lower upstream target is far more achievable even on longer copper phone lines. Lower usage of the upstream direction may also play a role here, but data is not available to prove or disprove this theory.

Cable services delivered 98.5% of advertised on average, with some countries achieving slightly over 100%. xDSL and FTTx services fared well too, providing 84.5% and 86.4% of advertised respectively.




Figure EU.2-7: Actual Upload Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, split by country




Figure EU.2-8: Actual Upload Speed of cable technology as a Percentage of Advertised Speed during peak periods, split by country




Figure EU.2-9: Actual Upload Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, split by country

It is interesting to note that Latvia and Lithuania – two of the strongest performing FTTx countries in the download metric – fared considerably poorer in this metric. The UK again falls below average across all technologies, but only slightly.

Exploring the actual upload figures delivered by each country and access technology again produces a very different picture of performance. Figures EU.2-10, EU.2-11 and EU.2-12 demonstrate this. Here we observe that once again the Nordic and Eastern European states deliver the best absolute figures, thanks to their deployment of FTTH services.

Latvia and Lithuania in particular, which had been identified earlier for delivering a relatively low percentage of advertised upload throughput, in fact delivered the highest absolute upload throughput of all countries studied.

The scale of the asymmetry between download and upload rates is most clearly seen in the xDSL figures. With an average of just 0.71Mbps for upload and 7.23Mbps for download, we have over a 10:1 ratio. FTTx services delivered just over a 2:1 average, demonstrating that newer technologies are moving towards more symmetric services.




Figure EU.2-10: Actual Upload Speed of xDSL technology during peak periods, split by country




Figure EU.2-11: Actual Upload Speed of cable technology during peak periods, split by country




Figure EU.2-12: Actual Upload Speed of FTTx technology during peak periods, split by country

D.2.3 Latency

Latency is an important metric that is often overlooked by consumers. Latency (or more precisely, round-trip latency) is the measure how long a single packet of data takes to go from point A to point B and back again. In this study, round-trip latency is measured between the panellists' homes and the nearest measurement server.

Given that every communication with Internet services involves the transmission and receiving of packets of data, latency affects everything we do on the Internet. It is particularly important for time-sensitive applications such as online gaming, video streaming and voice communications. The lower the latency, the more responsive the connection will be.

Different levels of latency are not a feature advertised with consumer broadband products, so it is impossible to compare against advertised levels. The access technology being employed by the ISP will often be the dominating factor with respect to latency.

Figures EU.2-13 through EU.2-15 depict the average round-trip latency per technology and country. On average, FTTx and cable services delivered very similar latencies. However, in some countries – such as Latvia, the Netherlands and Romania – latency for FTTx services was noticeably lower. This will be due to the prevalence of FTTH deployments in these countries, which do not have to use a xDSL-based last mile technology that incurs a significant latency overhead.

Average xDSL latencies were almost double those of Cable and FTTx services. With measurement servers not running in Cyprus, Malta and Iceland, it is not surprising that these islands see noticeably higher figures than other countries (as a key factor in latency is distance). Most surprising is the result for Spain, which exhibits high latency even though there are multiple measurement servers located in that country.




Figure EU.2-13: Latency of xDSL technology during peak periods, split by country




Figure EU.2-14: Latency of cable technology during peak periods, split by country




Figure EU.2-15: Latency of FTTx technology during peak periods, split by country

D.2.4 Packet Loss

Packet loss, like latency, is often overlooked when consumers are considering broadband products. This metric describes how many packets are lost as the traverse the path between the home and the computer they are communicating with. When packets are lost, the two parties involved in the communication will typically retransmit data in order to account for the loss. This takes time and when it reaches a certain level it will become very noticeable to users.

Realtime applications such as online gaming, video streaming and voice will again suffer the most from high packet loss.

Figures EU.2-16 to EU.2-18 below show packet loss figures during peak hours for each country considered in this study by technology. Most countries exhibit very low packet loss figures, never exceeding 1%, with only a few exceptions. All of the exceptions are seen on xDSL, which has double the average packet loss of the other two technologies studied. Higher loss on xDSL is normal and to be expected, as they typically use older copper lines which are more likely to suffer physical faults and defects.

Italy, Malta and Portugal stand out as the noticeable exceptions for delivering high levels of loss. Italy in particular is surprising, as multiple measurement servers were located in major cities in Italy, so the path between the user's home and the server should not have been that far.




Figure EU.2-16: Packet loss of xDSL technology during peak periods, split by country




Figure EU.2-17: Packet loss of cable technology during peak periods, split by country




Figure EU.2-18: Packet loss of FTTx technology during peak periods, split by country

Most countries exhibit packet loss figures that are between 0.2% and 0.5%, as seen in the figure below, with only three countries rising above 1% packet loss. It can be seen that Belgium, whose sample included only FTTx and cable services in this study, achieved the lowest overall figure. Other fibre-dominated countries such as Latvia also fared very well.

D.2.5 DNS Resolution Time and Failure Rate

DNS is a critical Internet service that allows you to turn hostnames (such as www.youtube.com) into IP addresses that your computer can communicate with. DNS services are typically provided by the ISP in order to provide a fast, nearby service for their users to use.

However, a poorly performing DNS service can lead users to perceive noticeable delays. This is particularly noticeable when browsing the web, which is heavily reliant upon DNS.

In theory, a good DNS deployment should provide DNS resolution time and failure rates better than or equal to the latency and packet loss figures. This is because the DNS servers are typically hosted inside the ISP's networks and therefore this traffic does not need to leave the ISP's network.

Figures EU.2-19, EU.2-20 and EU.2-21 show DNS resolution time for each country, split by technology. Again, the countries dominated by FTTH deployments deliver the best DNS resolution times, frequently at under 10 milliseconds. Belgium is a notable exception, with its FTTx (VDSL in this instance) services averaging far higher DNS resolution times than their latency figures.

In other technologies we observe that the trend generally matches the latency figures (as was expected). Cable services in Spain were another interesting exception, producing DNS resolution times over 10ms greater than their round-trip latency figures.




Figure EU.2-19: DNS Resolution Time of xDSL technology during peak periods, split by country




Figure EU.2-20: DNS Resolution Time of cable technology during peak periods, split by country




Figure EU.2-21: DNS Resolution Time of FTTx technology during peak periods, split by country

DNS failure rate by technology for each country is represented in Figures EU.2-22 through EU.2-24 below. With the exception of Cyprus and the Netherlands, all countries exhibited very low failure rates below 1%. The high failure rate value for the Netherlands is particularly surprising given their strong performance in prior metrics, and certainly warrants further investigation down to the ISP level.




Figure EU.2-22: DNS Resolution Failure Rate of xDSL technology during peak periods, split by country




Figure EU.2-23: DNS Resolution Failure Rate of cable technology during peak periods, split by country




Figure EU.2-24: DNS Resolution Failure Rate of FTTx technology during peak periods, split by country

In general, we observe that FTTx and Cable services deliver remarkably similar low failure rates. xDSL is noticeably more variable, and the similarity here to the packet loss chart suggests that this is driven by the underlying access technology rather than the ISP's services.

D.2.6 Web Browsing Speeds

Figures EU.2-25 to EU.2-27 below depict web page loading times in each country during peak hours, split by access technology. The test was performed to the public-facing websites of Facebook, Google and YouTube, all of which geographically host their servers for optimal customer performance.

Most countries' loading times range between 1 and 2 seconds on xDSL, whilst Cable and and FTTx services typically deliver page content in under 1 second.

Countries that exhibited the higher absolute downstream throughputs earlier are seen to perform well here. However, it can be observed that web browsing performance does not increase proportionally with increased throughput. Latency plays a key factor here too, so it is unsurprising that Cyprus (which produced high round-trip latency results) also provided the slowest web browsing times. A similar situation can also be observed for Malta.

It is interesting to note that Spain, which had previously been shown to suffer with some occurrences of high latency, did not stand out from the pack here. This suggests that the connectivity that our test websites have to the Spanish ISPs may be better than our own measurement servers.

One avenue for future improvement here would be to increase the sample of websites that are tested against. This may be extended to include local government websites, which, in theory, should be hosted locally within the country in question.




Figure EU.2-25: Web Page Loading Times of xDSL technology during peak periods, split by country




Figure EU.2-26: Web Page Loading Times of cable technology during peak periods, split by country




Figure EU.2-27: Web Page Loading Times of FTTx technology during peak periods, split by country

VoIP Jitter

Jitter is an important metric for those engaging in realtime communications. It can effectively be thought of as the consistency of latency. A broadband connection that jumped around between 10ms and 20ms latency continuously would have a high jitter value, and this behaviour would be troublesome for some realtime applications. Therefore, a lower jitter is desirable.

In this study we separately report upon downstream jitter and upstream jitter. In a two-way communication (such as a phone call), both are of course important. But important technology differences mean that these figures can often be quite different.

Downstream jitter, shown in Figures EU.2-28 to EU.2-30 below, proved to be low in most countries, never reaching figures above 2ms. Downstream jitter was higher for countries using xDSL technology than for all others, typically exhibiting figures between 1ms and 2ms. Italy achieved the highest level of jitter along with other notable examples such as Ireland and Malta. Belgium, however, achieved the lowest downstream jitter for cable services and was also one of the lowest for FTTx.




Figure EU.2-28: Downstream VoIP Jitter of xDSL technology during peak periods, split by country




Figure EU.2-29: Downstream VoIP Jitter of cable technology during peak periods, split by country




Figure EU.2-30: Downstream VoIP Jitter of FTTx technology during peak periods, split by country

Figures EU.2-31 to EU.2-33 present the results of the upstream jitter measurements. Unlike in the downstream jitter measurements, Italy displayed figures similar to most other countries of just above 2ms. Slovenia, on the other hand, displayed the highest figure for upstream jitter among xDSL technology countries, followed by Spain, Germany and Hungary. Spain also delivered the highest upstream jitter for FTTx.

Cable services on average delivered far higher upstream jitter levels than their xDSL and FTTx counterparts. This may appear counterintuitive at first, as cable services have generally performed very well in all metrics until now. The reason for cable services exhibiting higher jitter (particularly for the upstream) is due to the fact that they are based upon the concept of TDMA (Time Division Multiple

Access). Effectively the modem's time is divided into slots, during which it can either send or receive data. So if the modem is busy whilst the user tries to send a packet then that packet will have to wait in a queue until there is an opportunity to send it. This can result in small but frequent variations in packet delays, which is effectively what jitter represents.

It is important to note that whilst upstream jitter is often noticeably higher for cable networks, its level is often so low that it would be unnoticeable for almost all use cases. For example, most Voice over IP (VoIP) phones have a de-jitter buffer of at least 25ms, meaning jitter under 25ms would not affect call quality at all.




Figure EU.2-31: Upstream VoIP Jitter of xDSL technology during peak periods, split by country




Figure EU.2-32: Upstream VoIP Jitter of cable technology during peak periods, split by country




Figure EU.2-33: Upstream VoIP Jitter of FTTx technology during peak periods, split by country

European Commission

Quality of Broadband Services in the EU

106 pages



