



Quality of Broadband Services in the EU

October 2014

FINAL REPORT

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

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



SamKnows Limited

Internal identification

Contract number: 30-CE-0392545/00-77

SMART number 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.

ISBN 978-92-79-49486-4

DOI 10.2759/770879

© European Union, 2013. 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	8
A.1	Background	8
A.1.1	Purpose of the study	8
A.1.2	About SamKnows	10
A.1.3	Overview of Methodology	11
A.2	Key Findings	13
A.2.1	Summary of unweighted data	13
A.2.2	Performance by Technology	14
A.2.3	Performance by Country	15
A.3	Next Steps	16
B	METHODOLOGY AND DEFINITIONS	17
B.1	Methodology	17
B.1.1	An Open and Transparent Test Methodology	17
B.1.2	Measurement Methodology	18
B.1.3	Fixed broadband hardware probes	22
B.1.4	Software test (to initially check volunteer information)	24
B.1.5	Overview of network test nodes	27
B.1.6	European Commission Study Test Nodes	28
B.1.7	Control Suite	30
B.1.8	Reporting Infrastructure	36
B.1.9	Sample Plan Methodology	41
B.1.10	Final Panel Composition	42
B.1.11	Use of an All Volunteer Panel	43
B.1.12	Sample Size and Volunteer Selection	43
B.1.13	Panellist Recruitment Methodology	45
B.1.14	Validation of Volunteers' Service Tier	46
B.1.15	Protection of Volunteers' Privacy	47
B.1.16	Data Analysis Methodology	48
B.1.17	Collation of Results and Outlier Control	48
B.1.18	Peak Hours Adjusted to Local Time	49
B.1.19	Congestion in the Home Not Measured	49
B.1.20	Latencies Attributable to Propagation Delay	49
B.1.21	Limiting Factors	49
B.2	Definitions	50
B.2.1	Technology Splits	50
B.2.2	Scenario Matrix	51
B.2.3	Data Dictionary	52
B.3	Further information regarding sampling methodology	55

B.3.1	Identifying the Test Variables	55
B.3.2	Subdivisions of the Population	55
B.3.3	Sample Parameters	56
B.3.4	Participant Recruitment	57
B.3.5	Statistical Analysis Considerations	57
B.3.6	Data Validation	58
B.3.7	Impact of Unaccounted Variables (IPTV)	59
C	EU LEVEL ANALYSIS	60
C.1	Key Performance Indicators	60
C.1.1	Download Speed	60
C.1.2	Actual Download Speed split by hour of day and technology	63
C.1.3	Upload Speed	64
C.1.4	Latency	68
C.1.5	Packet Loss	70
C.1.6	DNS Resolution and Failure Rate	73
C.1.7	Web Browsing	77
C.1.8	VOIP	80
C.2	Comparison with the United States	85
C.2.1	Download Speed	85
C.2.2	Upload Speed	85
C.2.3	Latency	86
C.2.4	Packet Loss	86
D	COMPARISON BETWEEN COUNTRIES	87
D.1	Key Performance Indicators	87
D.2	Download and Upload Speeds	87
D.2.1	Download	87
D.2.2	Upload	91
D.2.3	Latency	96
D.2.4	Packet Loss	99
D.2.5	DNS Resolution Time and Failure Rate	102
D.2.6	Web Browsing Speeds	106
D.2.7	VoIP Jitter	109
E	EU LEVEL ANALYSIS - WEIGHTED	115
E.1	Key Performance Indicators	115
E.1.1	Download Speed	115
E.1.2	Actual Download Speed split by hour of day and technology	118
E.1.3	Upload Speed	119
E.1.4	Latency	122

E.1.5	Packet Loss	124
E.1.6	DNS Resolution and Failure Rate	127
E.1.7	Web Browsing	131
E.1.8	VOIP	134
E.2	Comparison with the United States	139
E.2.1	Download Speed	139
E.2.2	Upload Speed	139
E.2.3	Latency	140
E.2.4	Packet Loss	140
F	COMPARISON BETWEEN COUNTRIES - WEIGHTED	141
F.1	Key Performance Indicators	141
F.2	Download and Upload Speeds	141
F.2.1	Download	141
F.2.2	Upload	145
F.2.3	Latency	150
F.2.4	Packet Loss	153
F.2.5	DNS Resolution Time and Failure Rate	155
F.2.6	Web Browsing Speeds	159
F.2.7	VoIP Jitter	162
G	COMPARISON BETWEEN COUNTRIES - WEIGHTED	167
G.1	Key Performance Indicators	167
G.2	Download and Upload Speeds	167
G.2.1	Download	167
G.2.2	Upload	171
G.2.3	Latency	175
G.2.4	Packet Loss	177
G.2.5	DNS Resolution Time and Failure Rate	179
G.2.6	Web Browsing Speeds	182
G.2.7	VoIP Jitter	184

About This Document

This document is the third and final in a series of reports, which have been commissioned by the European Commission and completed by SamKnows over a three year period. For the purposes of this study 8,582 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 and the analysis of performance against key indicators across the EU.

In the previous two reports produced by SamKnows, all results were unweighted, however in this report we present both weighted and un-weighted data based on the number of subscribers per ISP and per technology within each country.

NRAs at all 30 countries included in this study provided SamKnows with market share information which allowed the inclusion of weighted results, and more information can be found within sections D and E of this report.

The analysis in this report is carried out on data collected in the month of October 2014. SamKnows continues to look for volunteers to participate in studies throughout Europe 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

A Executive Summary

A.1 Background

A.1.1 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 coverage and 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 fixed broadband coverage, with 97.2% of European households able to access broadband at the end of 2013, the figure drops to 89.8% in rural areas, and in some countries broadband covers just 63.5% or less of the rural households. In terms of take-up, there were 31 fixed broadband lines per 100 European citizens in June 2014.

The European Commission therefore commissioned a study on broadband performance to obtain reliable and accurate statistics of fixed 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 second in a series of three 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 and 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 December 2014 that average broadband speeds typically met advertised rates for in-country traffic with international performance proving almost equal if not slightly below advertised speeds for most major ISPs.

In the UK, SamKnows and Ofcom (the UK communications regulator) has found in November 2014 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.

Following the publication of the first and second European report we received considerable feedback from multiple stakeholders, including national regulators from member states, ISPs, content providers, academics and the public. Working with the European Commission, we have reviewed this feedback and incorporated changes where appropriate.

It is the hope of SamKnows that this third report will continue to 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.

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 8.27Mbps, 66.57Mbps and 53.09Mbps respectively, although there was much variation between countries. All technologies display a noticeable improvement in performance since the previous year, particularly cable technology which exhibits a 27.5% increase. 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 third and final report in the SamKnows/European Commission study, SamKnows have received feedback from the panellists that they would like to continue collecting data. Therefore 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 and local regulators. Any follow up to this project is done under the full responsibility of SamKnows exclusively without any contractual relationship with the European Commission.

A.1.3 **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 28 Member States of the European Union, as well as 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 measure the actual achieved speed. The study focuses on wired (fixed) technologies, specifically xDSL, Cable and FTTx and excludes wireless (mobile). The ISPs studied are those considered to be the two or three largest by subscribers in their national markets 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 a national level. 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 a minimum of 40 panellists per technology are required in order to be statistically valid to report upon. These panellists must be representative of the broadband consumers in each country. More details on this can be found in section B.3.3.

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 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
- Video streaming

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, further to feedback from national Regulatory Authorities after the first and second report, that enhancements can be made to the sample plan and data processing. Based on this, all the Regulatory Authorities provided SamKnows with market share information. This was used to ensure the sample in each country was accurate, and if allow weights to be used in the analysis.

A.2

Key Findings

A.2.1

Summary of unweighted data

This study presents the results of measurements taken from 8,582 measurement devices in October 2014. These devices were spread across 30 countries, 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 results included in this summary are based on unweighted data, however weighted results can be seen in sections E and F of this report.

The average download speed across all countries was 38.19Mbps during peak hours, a 25.7% increase from the previous year, slightly lower than the 39.69Mbps observed during the 24-hour measurement period. Average download speeds have therefore increased by nearly 10Mbps since October 2013, when the figures were 30.37Mbps and 31.72Mbps during the peak and 24-hour periods respectively. Peak period download speed represents 75.9% of the advertised headline speed, slightly up the 75.6% figure of 2013. 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 or by technology.

Whilst this compares poorly to the USA's average of 101% 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, particularly cable technology. xDSL services averaged 8.27Mbps in Europe and 7.67Mbps in the US. Cable services averaged 66.57Mbps in Europe and 25.48Mbps in the US. The same pattern was found for FTTx services too, with Europe averaging 53.09Mbps and US achieving 41.35Mbps.

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 10.94Mbps, increasing by 35.6% from 8.07Mbps since the last testing period. This figure represents 90.4% 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 speed was broadly similar across all countries and technologies, with most 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.

A.2.2 Performance by Technology

The headline figure of 78.3% of advertised download speed masks many important 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 86.5% and 83.1% respectively. Note that FTTx includes not only FTTH but also VDSL.

All technologies saw a small decrease 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.

Cable services achieved the fastest speeds in absolute terms, at 66.57Mbps. FTTx services achieved 53.09Mbps, whilst xDSL services lagged far behind at 8.27Mbps on average. This is a change from October 2013 as all technologies see a major improvement in actual performance in spite of generally exhibiting slightly higher throughput as a percentage of advertised rates.

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 with 25.23Mbps. 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 more modest 8.33Mbps and 0.85Mbps in comparison. Like with download speed, all technologies exhibit a small improvement in average performance since October 2013.

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 27.01ms, slightly lower than in October 2013. This figure is largely dictated by the technology in use, with xDSL averaging 37.36ms and cable and FTTx averaging 19.22ms and 20.16ms 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 0.32% across Europe at peak times, also an improvement from the previous testing period. Within this the figure was a considerably higher packet loss for xDSL services with 0.45% than cable packet loss of 0.23% and FTTx packet loss of 0.21%, which is a considerable improvement from 0.39% in October 2013. Interestingly, packet loss for xDSL was sharply lower at 0.30% in the 24-hour testing period. 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 Country

Download performance varied considerably by country. This situation is primarily driven by the technologies that have historically been deployed in those countries. Once again, the poorly performing states 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 20Mbps or 22Mbps, as was the case in the previous measurement period.

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 advertised speeds, this is not universally the case. Countries such as Poland and Slovakia are good examples of countries providing desirable xDSL services, achieving 85.6% and 86% of advertised speeds respectively. This may suggest 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.

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

Next Steps

This document represents the third and final in a series of three studies into the performance of broadband services across Europe.

SamKnows have received feedback from panellists that they would like to continue to collect data, and therefore the project will remain live and collecting data. This data can then be made available to the local regulators and within specialist reports published by SamKnows. Any follow up to this project is done under the full responsibility of SamKnows exclusively without any contractual relationship with the European Commission.

Following the publication of this report, we expect to continue receiving considerable feedback from multiple stakeholders, including national regulators from member states, ISPs, content providers, academics and the public.

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 continue building a collaborative working group with ISPs from across the 30 countries. 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, panellists 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 however has not been implemented in this report.

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.1.3 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

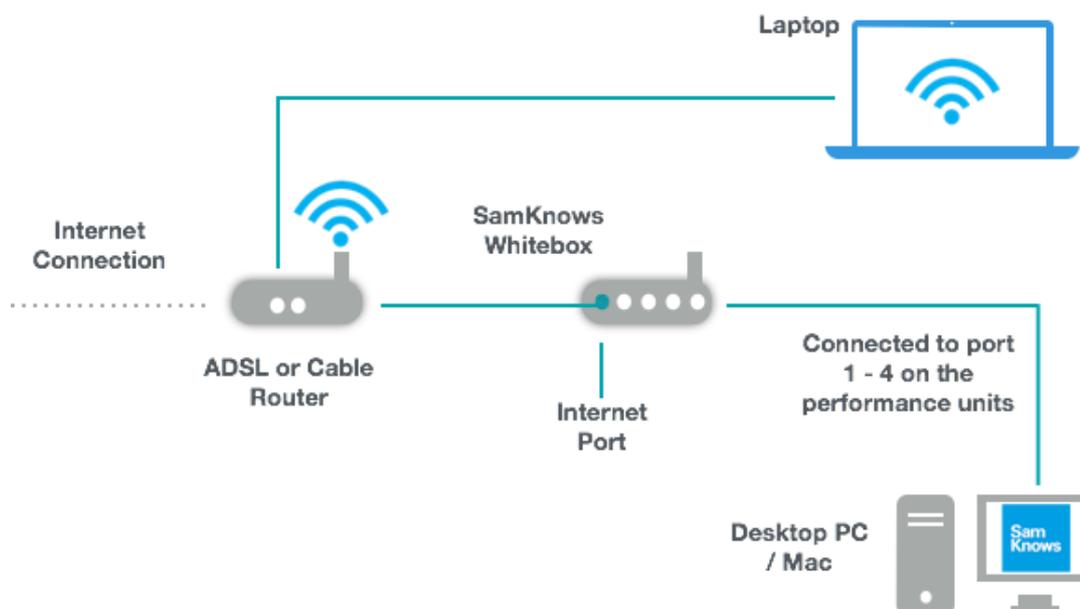
Whitebox 3.0 (SK-TL-WDR3600)

The SK-TL-WDR3600 can accurately measure fixed-line broadband connections of up to 500Mb/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
- 2.4GHz and 5GHz wireless, 2 antennas
- Single DC power (12V @ 1500mA)
- Power draw: 5W
- Weight: 962g

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, various heuristics are used to determine the correct wireless SSID to monitor, such as looking for an SSID with a MAC adjacent to that of the wired default gateway. If all of these heuristics fail, then the Whitebox falls back to selecting the SSID with the strongest signal. Only traffic from this SSID is monitored. In Whiteboxes that have 5GHz radios, the approach described here is used on both the 2.4GHz and 5GHz channels.

If the user's modem/router fully implements UPnP byte counters, then the probe will use this to monitor cross-traffic. Other techniques (such as HTTP or SNMP) may be employed, depending on the specific installation environment. However no tests are run when traffic is seen to be passing wirelessly.

B.1.4 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.1.5 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 second 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 only a small number of ISPs have 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.1.6 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. There were 35 test servers for the March 2012 measurement which was increased to 61 for the October 2013 measurement and finally 71 for the October 2014 measurement with a much wider distribution across Europe.

Measurement servers were located as follows:

Location	Hostnames
Amsterdam, Netherlands	ispmon.samknows.mlab1v4.ams01.measurement-lab.org ispmon.samknows.mlab2v4.ams01.measurement-lab.org ispmon.samknows.mlab3v4.ams01.measurement-lab.org ispmon.samknows.mlab1v4.ams02.measurement-lab.org ispmon.samknows.mlab2v4.ams02.measurement-lab.org ispmon.samknows.mlab3v4.ams02.measurement-lab.org n1-amsterdam-nl.samknows.com
Athens, Greece	ispmon.samknows.mlab2v4.ath01.measurement-lab.org ispmon.samknows.mlab3v4.ath01.measurement-lab.org ispmon.samknows.mlab1v4.ath02.measurement-lab.org ispmon.samknows.mlab2v4.ath02.measurement-lab.org ispmon.samknows.mlab3v4.ath02.measurement-lab.org
Hamburg, Germany	ispmon.samknows.mlab1v4.ham01.measurement-lab.org ispmon.samknows.mlab2v4.ham01.measurement-lab.org ispmon.samknows.mlab3v4.ham01.measurement-lab.org
London, UK	ispmon.samknows.mlab1v4.lhr01.measurement-lab.org ispmon.samknows.mlab2v4.lhr01.measurement-lab.org ispmon.samknows.mlab3v4.lhr01.measurement-lab.org n1-the1.samknows.com n2-the1.samknows.com

	n3-the1.samknows.com n4-the1.samknows.com n5-the1.samknows.com n6-the1.samknows.com
Leeds, UK	iispmon.samknows.mlab1v4.lba01.measurement-lab.org ispmon.samknows.mlab2v4.lba01.measurement-lab.org ispmon.samknows.mlab3v4.lba01.measurement-lab.org
Madrid, Spain	ispmon.samknows.mlab1v4.mad01.measurement-lab.org ispmon.samknows.mlab2v4.mad01.measurement-lab.org ispmon.samknows.mlab3v4.mad01.measurement-lab.org
Milan, Italy	ispmon.samknows.mlab2v4.mil01.measurement-lab.org ispmon.samknows.mlab3v4.mil01.measurement-lab.org
Turin, Italy	iispmon.samknows.mlab1v4.trn01.measurement-lab.org ispmon.samknows.mlab2v4.trn01.measurement-lab.org ispmon.samknows.mlab3v4.trn01.measurement-lab.org
Paris, France	ispmon.samknows.mlab1v4.par01.measurement-lab.org ispmon.samknows.mlab2v4.par01.measurement-lab.org ispmon.samknows.mlab3v4.par01.measurement-lab.org n1-paris-fr.samknows.com
Stockholm, Sweden	ispmon.samknows.mlab1v4.arn01.measurement-lab.org ispmon.samknows.mlab2v4.arn01.measurement-lab.org ispmon.samknows.mlab3v4.arn01.measurement-lab.org
Hudiksvall, Sweden	n1-hudiksvall-se.samknows.com
Dublin, Ireland	ispmon.samknows.mlab1v4.dub01.measurement-lab.org ispmon.samknows.mlab2v4.dub01.measurement-lab.org ispmon.samknows.mlab3v4.dub01.measurement-lab.org
Ljubljana, Slovenia	ispmon.samknows.mlab1v4.lju01.measurement-lab.org ispmon.samknows.mlab2v4.lju01.measurement-lab.org ispmon.samknows.mlab3v4.lju01.measurement-lab.org
Cyprus	ispmon.samknows.mlab1v4.lca01.measurement-lab.org ispmon.samknows.mlab3v4.lca01.measurement-lab.org
Prague, Czech Republic	ispmon.samknows.mlab1v4.prg01.measurement-lab.org ispmon.samknows.mlab2v4.prg01.measurement-lab.org ispmon.samknows.mlab3v4.prg01.measurement-lab.org ispmon.samknows.mlab4v4.prg01.measurement-lab.org
Stavanger, Norway	ispmon.samknows.mlab1v4.svg01.measurement-lab.org ispmon.samknows.mlab2v4.svg01.measurement-lab.org ispmon.samknows.mlab3v4.svg01.measurement-lab.org
Vienna, Austria	ispmon.samknows.mlab1v4.vie01.measurement-lab.org ispmon.samknows.mlab2v4.vie01.measurement-lab.org ispmon.samknows.mlab3v4.vie01.measurement-lab.org
Belgrade, Serbia	ispmon.samknows.mlab1v4.beg01.measurement-lab.org ispmon.samknows.mlab2v4.beg01.measurement-lab.org ispmon.samknows.mlab3v4.beg01.measurement-lab.org
Bucharest, Romania	n1-bucharest-ro.samknows.com
Warsaw, Poland	n1-warsaw-pl.samknows.com
Riga, Latvia	n1-riga-lv.samknows.com
Luxembourg	n1-ept-lu.samknows.com

Sofia, Bulgaria	n1-sofia-bg.samknows.com
Budapest, Hungary	n1-budapest-hu.samknows.com
Zagreb, Croatia	n1-unizg-zagreb-hr.samknows.com
Warsaw, Poland	n1-warsaw-pl.samknows.com

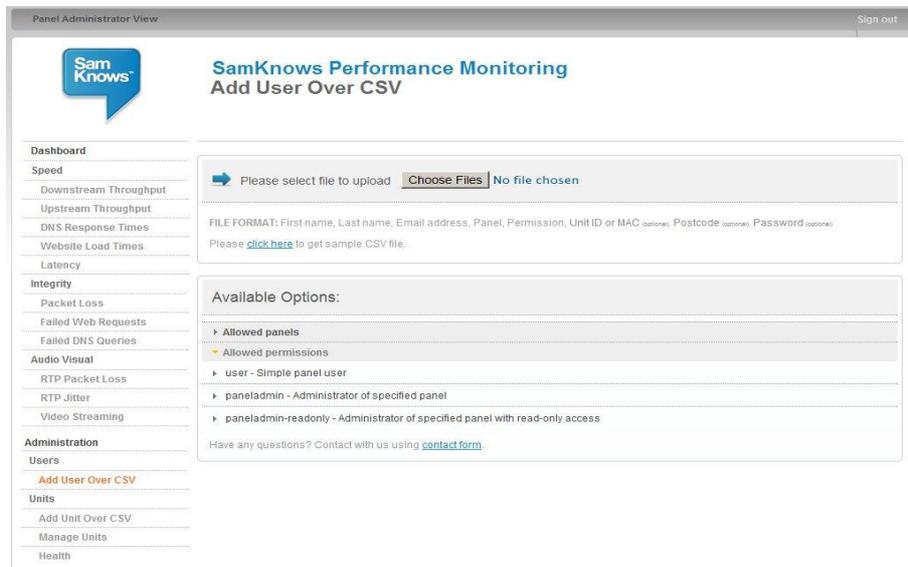
B.1.7 Control Suite

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	Once 12am-6am	<30Mbps =	54MB
		Once 6am-12pm	6MB	
		Once 12pm-6pm	30-50Mbps =	108MB
		Every hour 6pm-12am	12MB file size	
			>50Mbps =	>~540MB
			10 seconds duration	
Upload speed	1 off-net test node	Once 12am-6pm	<10Mbps =	18MB
		Once 6am-12pm	3MB fixed	
		Once 12pm-6pm	size	
		Once 6pm-12pm	10-20Mbps =	36MB
			6MB	
			>20Mbps =	>~216MB
			10 seconds duration	
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

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.

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 = Total Throughput of the line = 197469.668017 + 156168.655454
+ 161044.475879 = 514682 (Bps) * 0.000008 = 4.12 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.

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 including FCC (United States of America), Ofcom (UK), IDA (Singapore), CRTC (Canada) and Anatel (Brazil).

B.1.8 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.

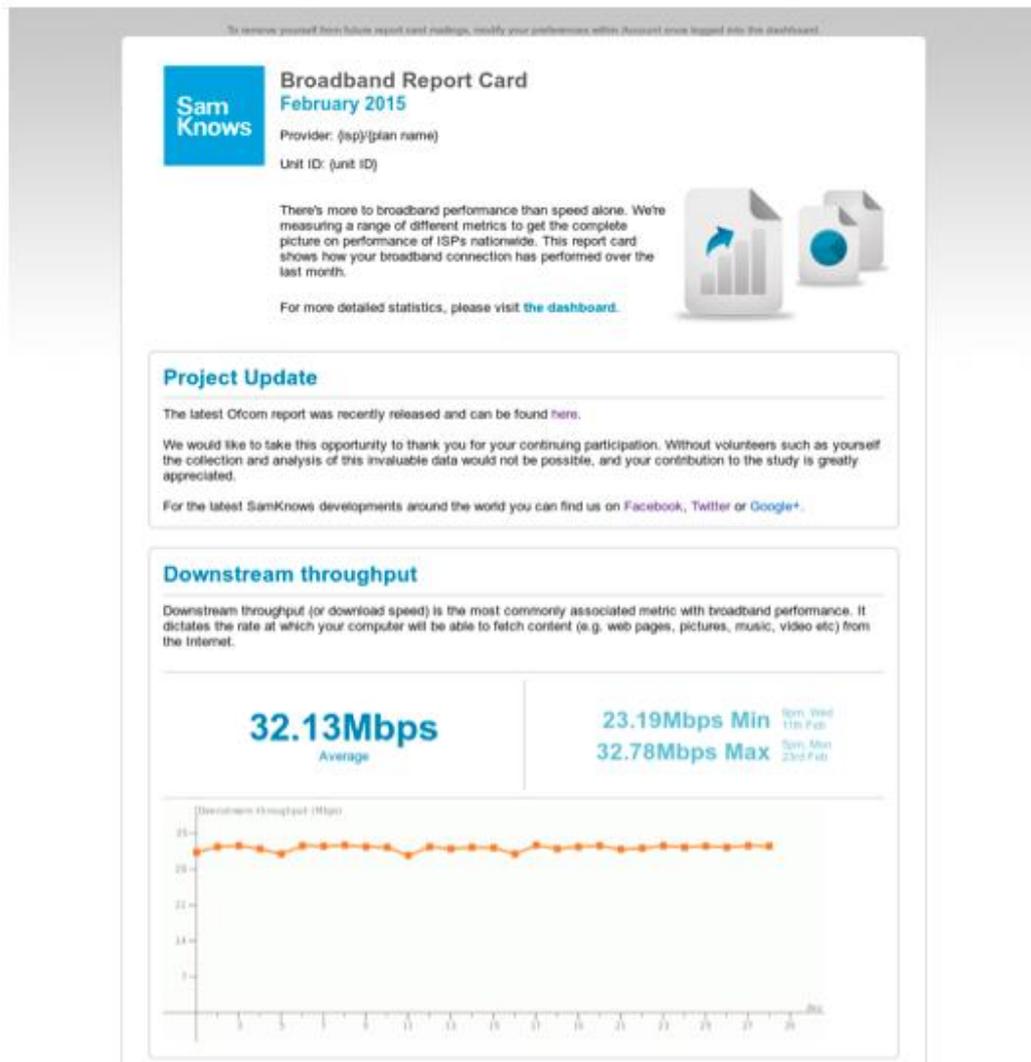
Performance graphs and statistical results are collected into the areas of speed, integrity and multimedia performance and are viewable both in dashboard overview format and in extended detail in each relevant section.

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 complement 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



B.1.9 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 the following European Commission's multi annual studies¹: "Broadband Coverage in Europe", "Broadband Internet Access Cost (BIAC)" and COCOM report on fixed broadband access lines, as well as from commercial sources. The data in these studies are broken down according to country, technology and speed tier. 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.

Within each country and technology type, we then analyse the data taken from the 2012 European Commission documents 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, in 2010 the proportion of broadband users in Belgium with cable technology and a speed tier of 0-2Mbps was 0%, and those with a tier of 2-8Mbps was 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 representative 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 2014 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.

Note: Because of the ever-changing nature of broadband services, the panel composition will change over time to remain consistent with the overall

¹ <http://ec.europa.eu/digital-agenda/en/fast-and-ultra-fast-internet-access-analysis-and-data>

market. However the sample plan methodology will remain consistent throughout.

B.1.10 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 panellists, which contains both current and legacy products across all countries, is therefore displayed below, with sufficient sample sizes being highlighted:

Country	xDSL Panel	Cable Panel	FTTx Panel
Austria	18	31	12
Belgium	21	322	108
Bulgaria	21	34	131
Croatia	25	12	1
Cyprus	28	16	0
Czech Republic	74	94	26
Denmark	112	36	203
Estonia	10	3	7
Finland	74	34	72
France	299	34	37
Germany	203	130	133
Greece	219	0	26
Hungary	44	121	12
Iceland	8	0	34
Ireland	44	55	24
Italy	533	0	47
Latvia	10	9	36
Lithuania	25	11	137
Luxembourg	15	10	27
Malta	19	48	7
Netherlands	100	291	140
Norway	32	18	79
Poland	158	127	54
Portugal	84	152	120
Romania	18	23	177
Slovakia	40	24	37
Slovenia	49	30	115
Spain	292	143	103
Sweden	34	23	136
United Kingdom	907	386	843
TOTAL	3516	2217	2849

B.1.11 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.1.12 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:

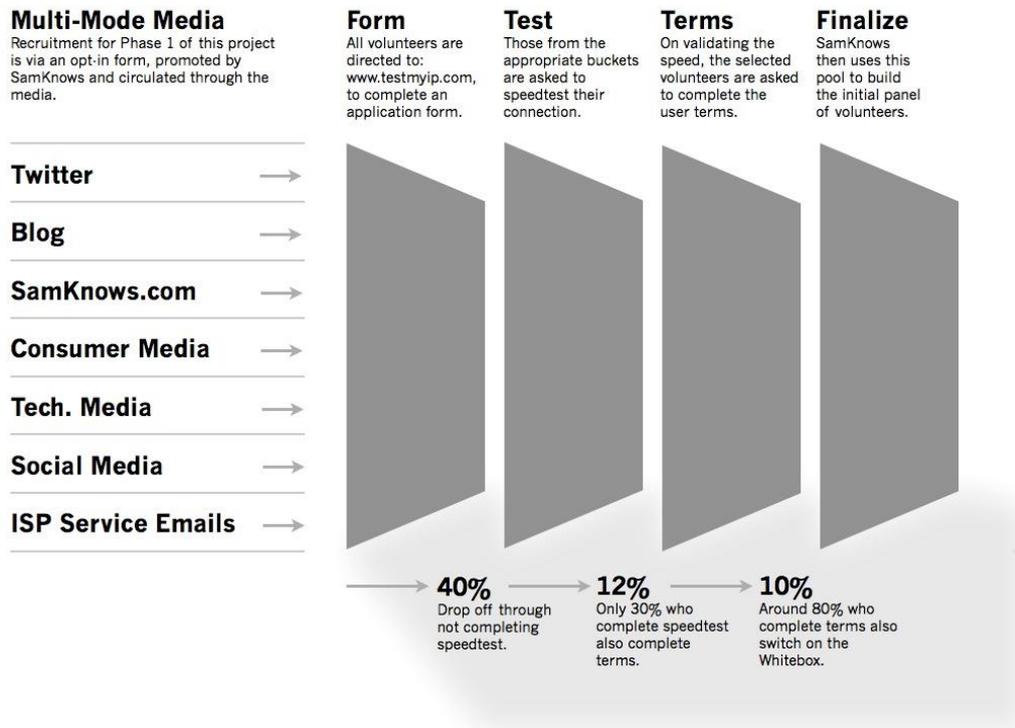
- The existing 3,500 panelists whom remained reporting from the first two studies in March 2012 and October 2013.
- The panel of EU broadband subscribers was drawn from a pool of over 250,000 volunteers following an ongoing recruitment campaign that ran from February 2011.
- The volunteer sample was organized with a goal of covering major ISPS in the 28 Member States + Iceland and Norway across the available broadband technologies: xDSL, Cable, FTTx.
- 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 continues.

- 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.

B.1.13 **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 of 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.1.14 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.

Note that without ISP participation, it is possible that some panelists may have been allocated an incorrect technology or service tier.

B.1.15 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.1.16 Data Analysis Methodology

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.1.17 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 40 reporting Whiteboxes.

The resulting final sample of data for October 2014 was 8,582 participants.

B.1.18 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.1.19 Congestion in the Home Not Measured

Download, upload, latency, and packet loss measurements were taken between the panellist's home gateway and the dedicated test nodes. Web browsing measurements were taken between the panellist'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.1.20 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.1.21 Limiting Factors

A total of 10,418,841,762 measurements were taken across 66,466,182 unique tests. All scheduled tests were run, aside from when monitoring units detected concurrent use of bandwidth.

B.2 Definitions

B.2.1 Technology Splits

Results in sections C, and D are often split by access technology. This report defines these technologies as ‘xDSL’, ‘Cable’ and ‘FTTx’.

Note that without ISP participation in validation of panelists, it is possible that some panelists may have been allocated an incorrect technology or service tier.

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 – which is the category in which some countries in Europe market the product). However it is not always possible to distinguish between xDSL and VDSL.

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_udplateness)	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 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.
Very high/unstable jitter (>100ms)	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

B.2.3 Data Dictionary

curr_dns.csv

unit_id	Unique identifier for an individual unit
dtime	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
dtime	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
dtime	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
tftb_total_time	Total duration of the time-to-first-byte summed together, if made sequentially
tftb_min_time	Shortest time-to-first-byte duration
tftb_avg_time	Average time-to-first-byte duration
tftb_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 these criteria. Please note that over the course 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. Within the last EC report, we used a minimum of 50 panelists to be consistent with our studies on behalf of Ofcom (UK). However following separate independent analysis on all SamKnows global speed test data, the conclusion is that 40 to 45 Whiteboxes (or greater) provides an acceptable level of sampling reliability, can be considered asymptotically normally distributed, and will be robust for statistical analysis. The full analysis can be found here:

<http://www.samknows.com/broadband/methodology>

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.

Please note that for the third 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. However, the European Commission study (2014) 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 worse the effective performance for these measures is likely to be wherein

download and upload will be lower and latency and packet loss will be higher. Not accounting for distance risks introducing an upward bias for these metrics. To illustrate, Malta, Spain and Portugal 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 measurement servers located within the country (Madrid), whereas Portugal and Malta do not which may be a reason for the higher latency. 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 a very high packet loss for xDSL technology and also FTTx technology. 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, Portugal's latency for cable technology is above average to a similar degree as Malta and Spain, 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. For example comparing all tests run on the line over the whole period, and specifically at peak time when IPTV is likely to be used against off-peak times. This is something that is being developed by SamKnows analysts, with the intention of this functionality being built-in to the user reporting for subsequent reports. 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 Analysis²

c.1 Key Performance Indicators

The data within sections C and D are unweighted as per the two previous EC reports produced by SamKnows.

c.1.1 Download Speed

Figure EU.1-1 and Fig EU.1-2 display actual download speed as a percentage of advertised speed over the 24-hour and peak periods, split by access technology. Cable technology outperforms all other access technology, achieving the highest level of download throughput as a percentage of advertised speed with 86.51% during the peak period, representing a slight drop from results over the 24-hour period. FTTx and xDSL technologies deliver 83.14% and 63.32% of advertised speed respectively during this period. Only FTTx technology displays a small improvement from 82.7% of advertised speed achieved in the last testing period of October 2013, although 24-hour results were slightly higher in the previous measurement period. All other technologies experienced a small decline in percentage terms during both the peak and 24-hour periods. However, this is due to an increase in average headline speeds of each access technology as opposed to a real decline in overall performance. All access technologies experienced significant improvements in actual throughput, especially cable technology.

Please note that the figures below are not obtained by dividing the average actual speed through the average advertised speed. They are computed on a per-panellist basis and averaged to form an overall figure. This approach is therefore a mean of ratios as opposed to a ratio of means.

	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Technology and Period								
Actual Speed (Mbps)								
Oct-14	8.27	8.5	66.57	70.04	53.09	54.65	38.19	39.69
Oct-13	8.13	8.31	52.21	54.59	47.74	50.06	30.37	31.72
Advertised Speed (Mbps)								
Oct-14	14.04	14.04	79.8	79.8	64.92	64.92	47.9	47.9
Oct-13	13.95	13.95	60.54	60.54	59.48	59.48	38.5	38.5
Actual/Advertised Speed								
Oct-14	63.32%	65.07%	86.51%	90.16%	83.14%	85.45%	75.89%	78.31%
Oct-13	63.80%	65.10%	89.50%	92.90%	82.70%	86.50%	75.60%	78.10%

² EU refers to the average of all the countries included in the sample, i.e. EU28 countries & Iceland and Norway

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

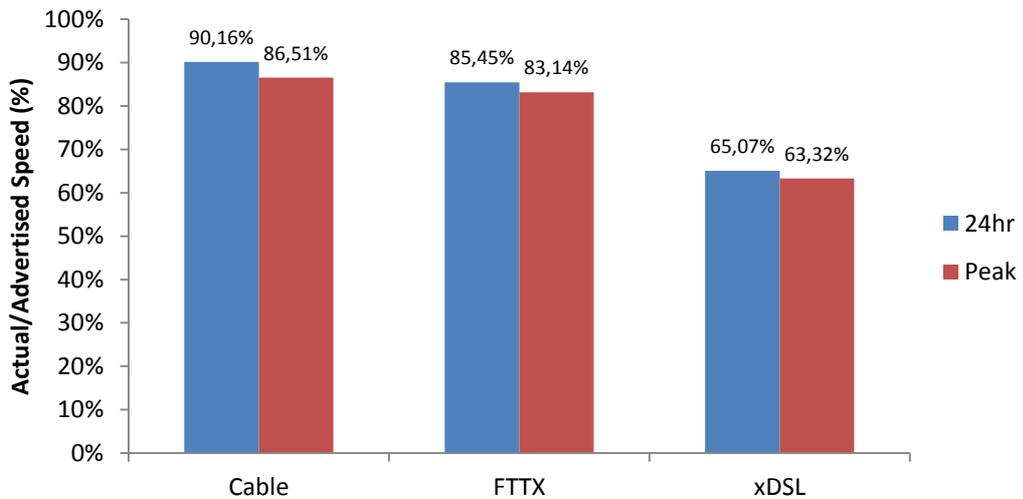


Figure EU.1-2: Actual Peak and 24-hour Period Download Speed as a Percentage of Advertised Speed, by technology (higher is better)

Figure EU.1-3 below shows actual download speed as a percentage of advertised speed, split by time of day and technology. All technologies exhibit a similar pattern throughout the day, with download speed declining slightly during the day, followed by a more significant drop in throughput during the peak period. xDSL exhibits the smallest amount of decline during peak hours compared to cable and FTTx technologies, with cable technology exhibiting the sharpest drop in throughput during this time. As was shown in figure EU.1-2, cable outperforms all other technologies.

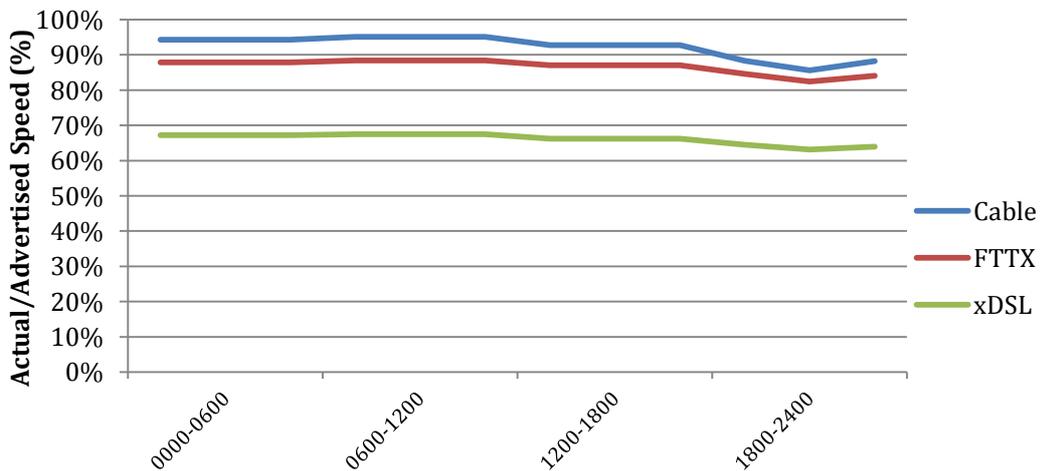


Figure EU.1-3: Actual Download Speed as a Percentage of Advertised Speed, by hour of day and technology (higher is better)

Figure EU.1-4 below shows actual download speed split by time of day and access technology. As seen in figure EU.1-3, all technologies display similar diurnal

patterns. Download speed experiences a small decline in actual throughput in the afternoon followed by a sharper drop in the peak period, as is indicated in figure EU.1-3. xDSL technology shows the smallest amount of change during this period, contrasting with cable technology which experiences the largest decline. Additionally, all technologies experienced significant improvements in actual download performance since October 2013, particularly cable technology.

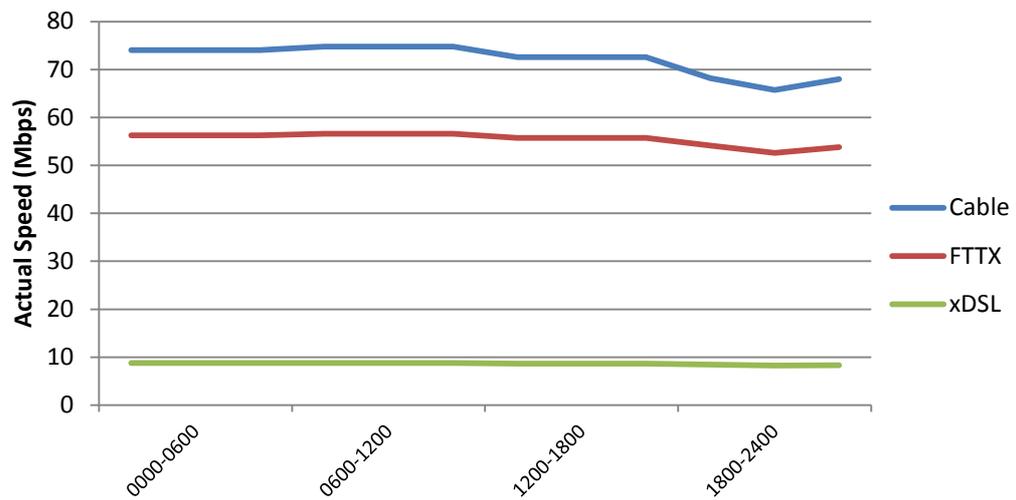


Figure EU.1-4: Actual Download Speed by hour of day and technology (higher is better)

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

Figure EU.1-5 presents a cumulative distribution chart of download speed as a percentage of advertised speed, split by access technology. This chart is meant to represent the percentage of consumers who receive at least a certain level of their

advertised broadband speed i.e. the chart gives an indication of what proportion of consumers receive their respective advertised speeds. Figures EU.1-2 to EU.1-4 focused on averages alone, but these can mask high levels of inconsistency. The cumulative distribution plot helps to show if there is a significant spread of results within the measurement samples.

For example, one technology may deliver 90% of advertised speed to all its users at all times, and another technology may deliver anything between 60% and 100% of advertised speed. Both may produce an average speed of 90%, thus proving hard to distinguish in charts EU.1-2 to EU.1-4 above. The cumulative distribution chart presents these differences clearly.

Figure EU.1-5 shows 80% of cable consumers receive at least 75% of advertised speed, FTTx consumers receive approximately 70% and xDSL consumers much less than other technologies with 36% of the advertised rate. In the case of cable technology, this appears to be a slight decrease from October 2013. This is due to higher overall advertised rates, with actual throughput seeing an improvement in performance. xDSL and FTTx technologies instead see a slightly greater percentage of consumers receiving better speeds. Results for xDSL consumers show a much wider distribution compared to cable and FTTx, with only 31% of xDSL consumers receiving 80% of advertised speed or better. It is also 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.x.

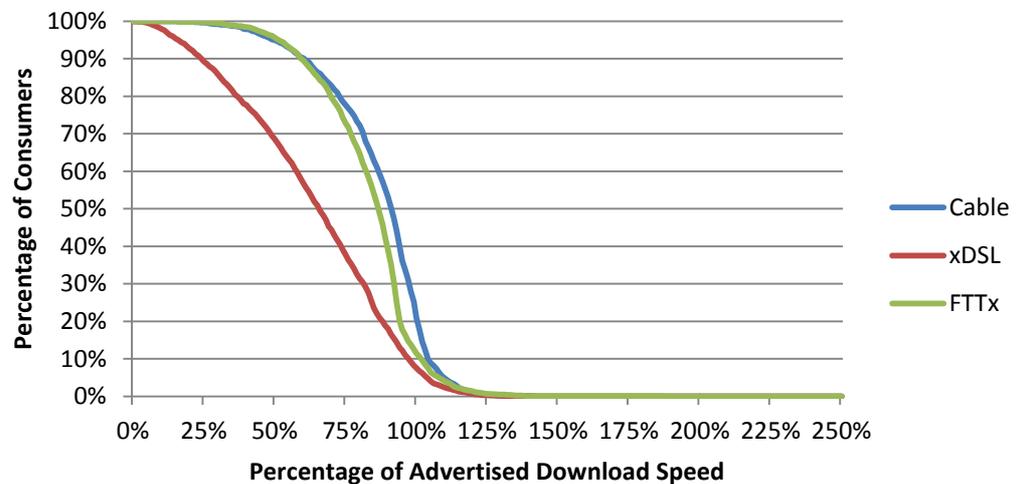


Figure EU.1-5: Cumulative Distribution of Download Speed as a Percentage of Advertised Speed, by technology

c.1.3 Upload Speed

Figures EU.1-6 and EU.1-7 display actual upload speed expressed as a percentage of advertised speed over the peak and 24-hour periods split by access technology. All technologies achieved a level of upload speed above 80% of advertised speed during both periods, with all technologies showing a small decline since October 2013. As was the case in the previous testing period, cable technology continues to

exceed the advertised speed during the peak and 24-hour periods despite the decline in performance in percentage terms, with almost no difference in throughput between both periods. Cable again outperforms all other technologies in percentage terms. All technologies show a very small decrease in upload throughput during the peak period.

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 computed on a per-panellist basis and averaged to form one overall figure, as opposed to dividing multiple averages together.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Actual Speed (Mbps)								
Oct-14	0.85	0.86	8.33	8.4	25.23	25.54	10.94	11.06
Oct-13	0.81	0.82	6.3	6.34	21.6	22.29	8.07	8.28
Advertised Speed (Mbps)								
Oct-14	1.05	1.04	8.43	8.43	27.91	27.91	11.94	11.94
Oct-13	0.96	0.96	6.25	6.25	24.84	24.84	9.05	9.05
Actual/Advertised Speed								
Oct-14	82.70%	83.19%	100.10%	100.64%	92.29%	93.09%	90.36%	90.97%
Oct-13	86.30%	86.70%	103.40%	104.00%	94.00%	96.30%	92.80%	93.80%

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

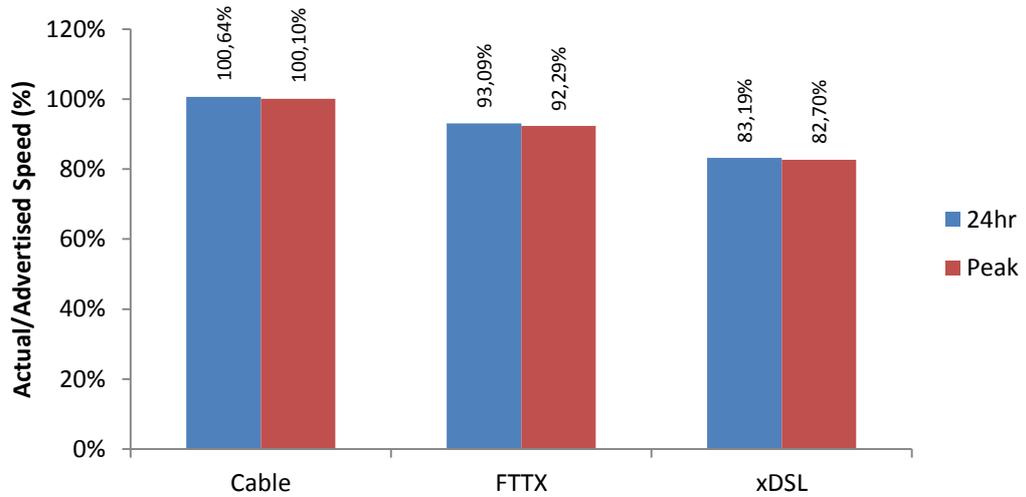


Figure EU.1-7: Actual Peak and 24-hour Period Upload Speed as a Percentage of Advertised Speed, by technology (higher is better)

Figure EU.1-8 shows actual upload speed as a percentage of advertised speed split by technology and time of day. As was the case in the previous measurement period, Cable and xDSL services deliver a very stable level of throughput during the day, experiencing a very small decline during peak hours. FTTX also displays a relatively stable level of throughput. All technologies show a slightly lower level of throughput in the later hours of the day.

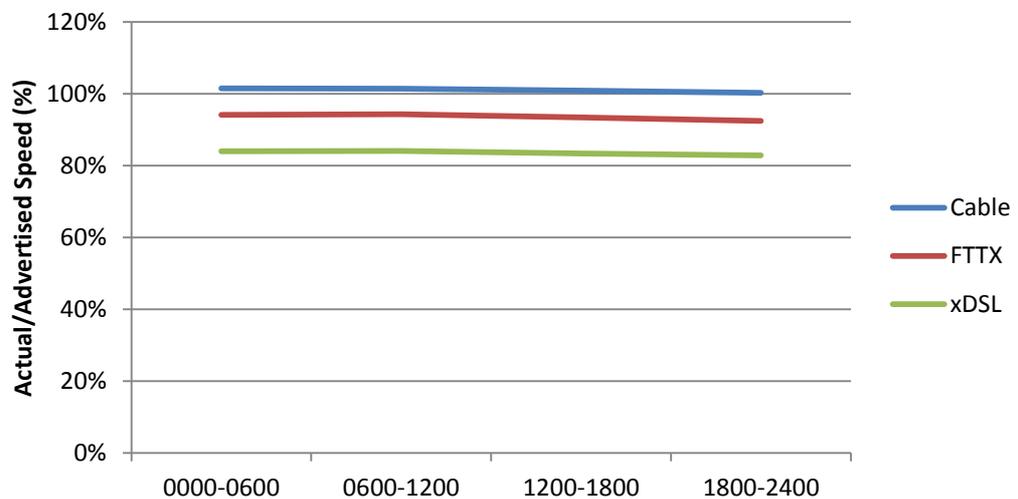


Figure EU.1-8: Actual Upload Speed as a Percentage of Advertised Speed, by hour of day and technology (higher is better)

Figure EU.1-9 shows actual upload speed by time of day and technology. The behaviour of throughput for all technologies is virtually identical to what is seen in EU.1-8 in the figure below. Cable and xDSL services deliver very stable throughput throughout the day. As was the case in October 2013, a decrease in throughput is more noticeable for FTTx technology throughout the afternoon and peak periods. Contrasting with EU.1-8, FTTx greatly outperforms all other technologies in real terms, indicating much higher advertised rates for FTTx based packages across the countries considered in this report.

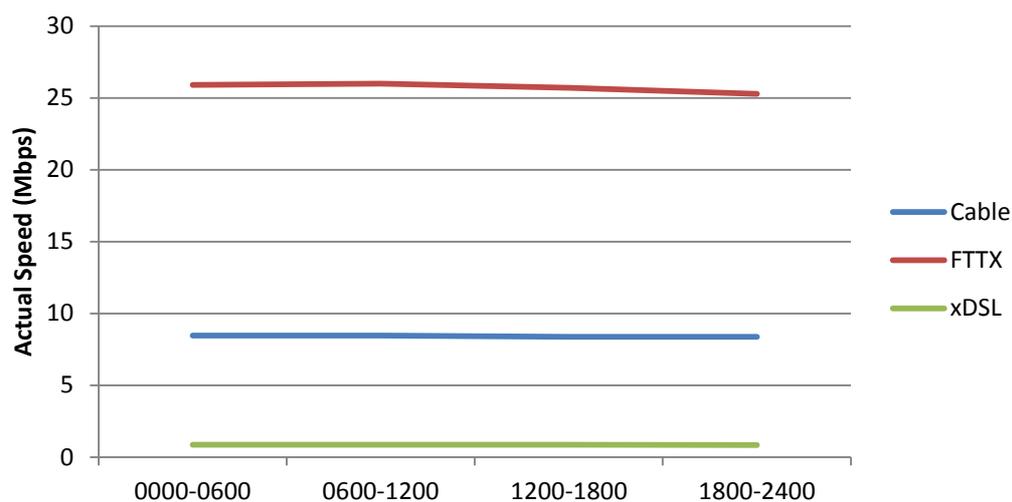


Figure EU.1-9: Actual Upload Speed by hour of day and technology (higher is better)

Figure EU.1-10 is a cumulative distribution chart for upload speed expressed as a percentage of advertised speed.

The distribution of results is very similar to the previous measurement period, with consumers generally receiving higher levels of advertised speed.

80% of cable consumers are receiving 95% of advertised speed. As was the case in October 2013, FTTx and xDSL are both spread over wider distributions compared to cable technology, with 80% of FTTx and xDSL consumers receiving 80% and 68% of advertised speed respectively. For FTTx technology, this is a significant improvement, in contrast to xDSL technology whose performance is virtually unchanged.

The distribution of results for xDSL consumers is much tighter for upload speed compared to download speed. This is caused by the asymmetric nature of broadband services, with upload speed provisioned at far lower rates.

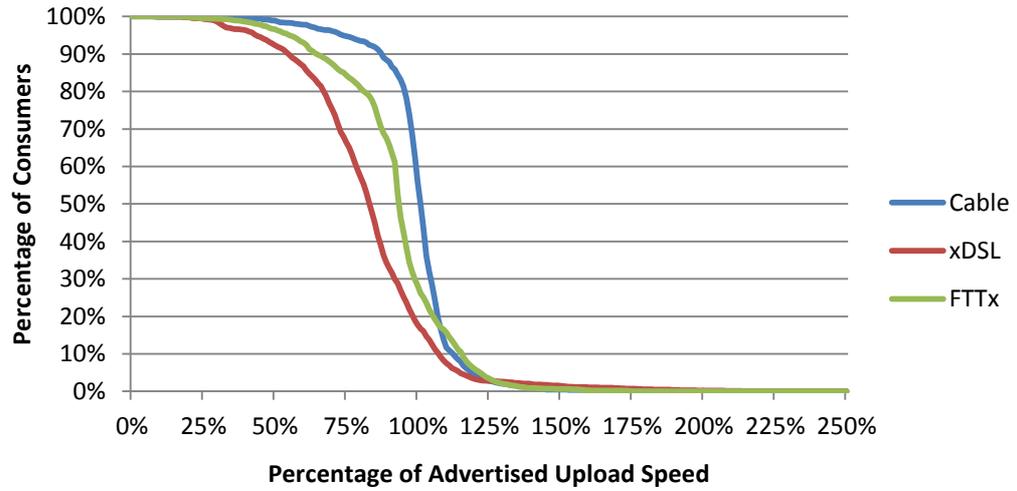


Figure EU.1-10: Cumulative Distribution of Upload Speed as a Percentage of Advertised Speed, by technology

C.1.4 Latency

Figures EU.1-11 and EU.1-12 show average latency over the peak and 24-hour periods across all types of access technology included in this report. xDSL and FTTx technologies display a slight increase in latency since October 2013, in contrast to cable technology which exhibits a slight improvement. Although most access technologies display a small increase in latency since October 2013, average latency across the EU is slightly lower. Peak period changes are also larger than they were in the previous testing period. Latency is lower during the 24-hour period compared to the peak period, with xDSL displaying the largest change in real terms.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Latency (ms)								
Oct-14	37.36	35.44	19.22	17.98	20.16	18.79	27.01	25.45
Oct-13	36.41	35.16	20.87	19.5	19.16	18.48	27.65	26.53

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

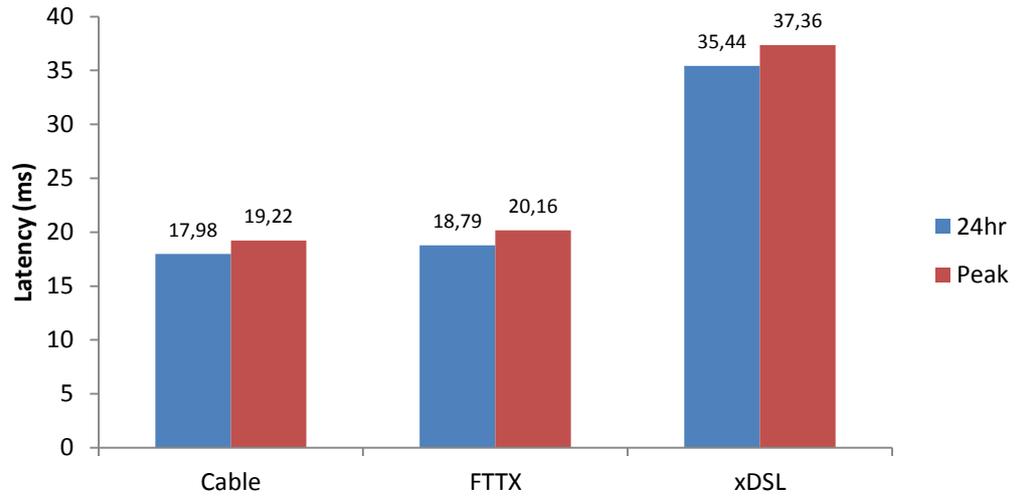


Figure EU.1-12: Peak period and 24-hour Latency, by technology (lower is better)

Figure EU.1-13 displays latency split by hour of day and technology. As mentioned above for figure EU.1-12, xDSL experiences the most noticeable increase in latency during the peak period. This increase is also more noticeable compared to the previous testing period. In contrast, FTTx latency has become more stable since October 2013.

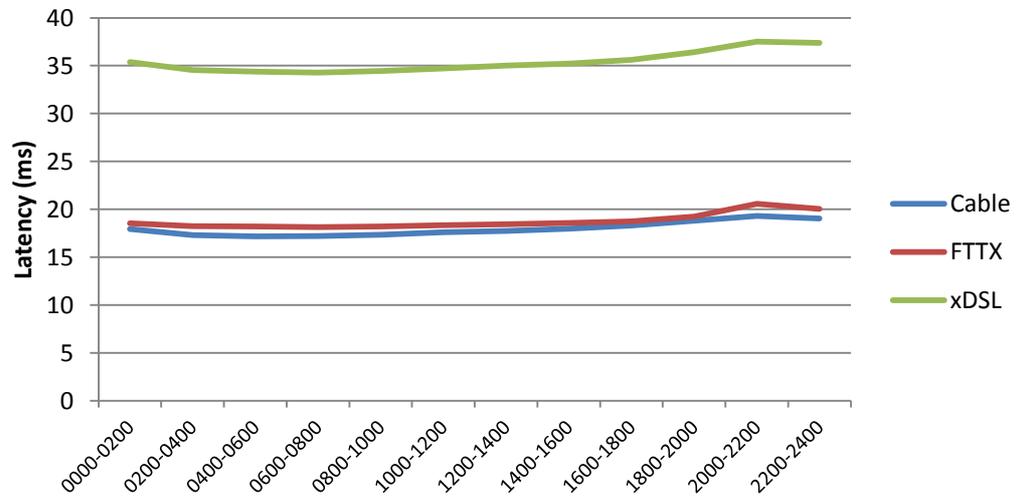


Figure EU.1-13: Latency by hour of day and technology (lower is better)

Figure EU.1-14 depicts the cumulative distribution of latency. This differs from the CDF charts shown previously for download and upload throughput as it does not display the percentage of advertised figures. This is because ISPs generally do not advertise latency for their broadband products (aside from exceptional cases), as latency can vary wildly depending on the host the user is communicating with. This CDF chart and all others that follow it will instead show the actual value unless stated otherwise.

Figure EU.1-14 shows cable and FTTx technologies have similar distributions of latency across its consumers, as was the case in October 2013. xDSL technology again exhibits a comparatively wider distribution, reflecting the behaviour of download throughput. The CDF plot below also shows that latency during this report's testing period is very close to what it was in October 2013, with 60% of cable and FTTx consumers experiencing latency of up to 17ms and 20ms respectively. It is expected behavior for cable and FTTx to experience lower latency compared to xDSL technology as they are less affected by the length of the "last-mile" cable than xDSL, with 60% of xDSL consumers experience 37ms or less. However, this is a significant improvement for xDSL technology, with 60% of consumers previously receiving 50ms or less.

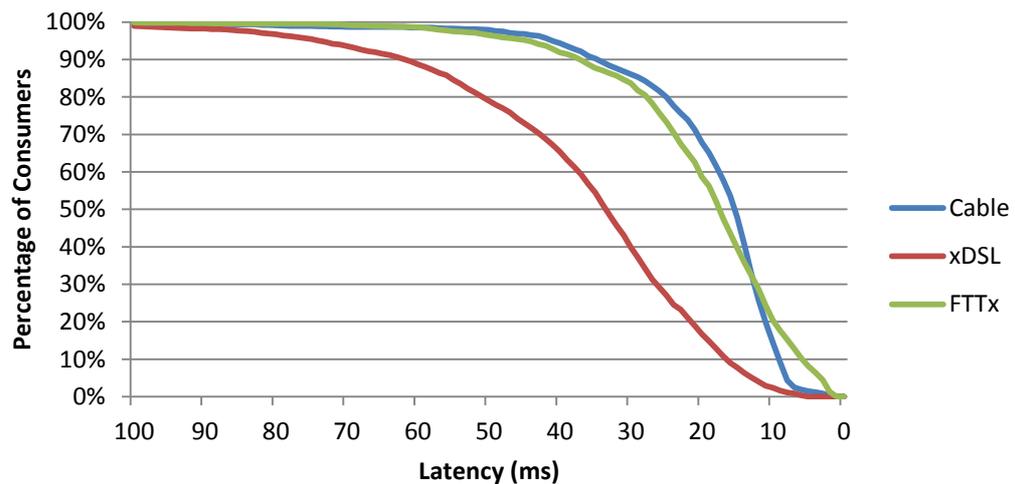


Figure EU.1-14: Cumulative Distribution of Latency, by technology

c.1.5 Packet Loss

Figure EU.1-15 and Figure EU.1-16 show the average packet loss over the peak and 24-hour periods split by access technology. xDSL exhibits the largest level of packet loss across both time periods compared to all other technologies as well as a significantly larger change between the peak and 24-hour periods. In contrast, cable and FTTx experience a very small increase in packet loss during peak hours. Packet loss for xDSL and FTTx technology is lower during this testing period compared to October 2013, with FTTx exhibiting a significant improvement.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Packet Loss (%)								
Oct-14	0.45%	0.30%	0.23%	0.20%	0.21%	0.17%	0.32%	0.23%
Oct-13	0.50%	0.35%	0.20%	0.18%	0.39%	0.22%	0.39%	0.27%

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

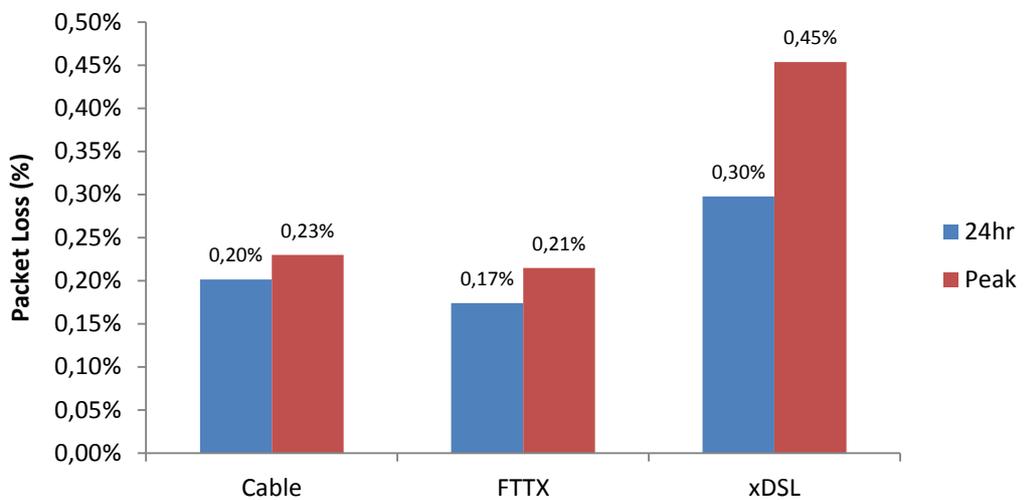


Figure EU.1-16: Peak period and 24-hour packet loss, by technology (lower is better)

Figure EU.1-17, which displays packet loss by hour of day and split by technology, demonstrates the noticeable increase in packet loss for xDSL services during the day, particularly in the peak period. All types of technology also behave very similarly throughout the day, slightly decreasing in the morning period followed by a sharp increase during the late afternoon and peak hours. These changes are more significant for packet loss of xDSL technology than for cable and FTTx. Packet loss for xDSL technology is also consistently higher than it is for cable and FTTx throughout the day.

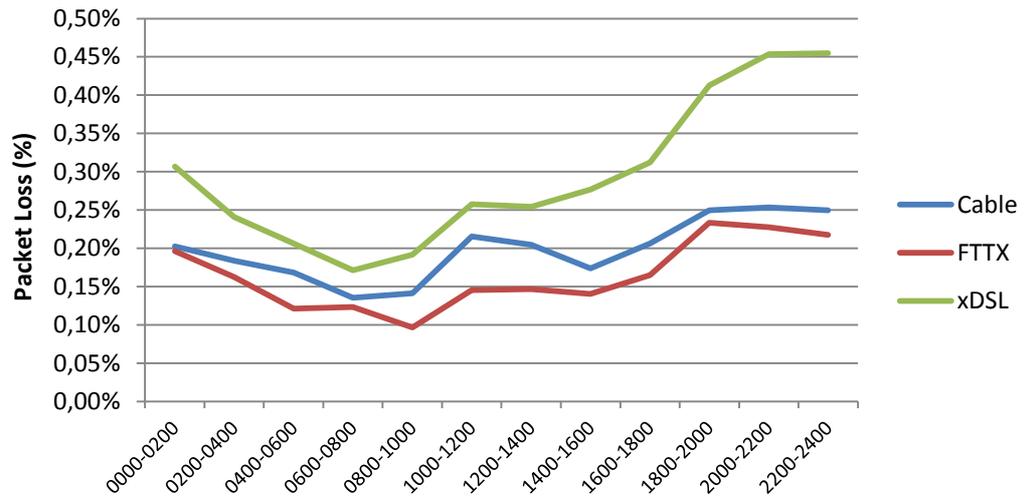


Figure EU.1-17: Packet Loss by hour of day and technology (lower is better)

Figure EU.1-18 is the cumulative distribution plot for packet loss. As was the case in October 2013, very few consumers experience high levels of packet loss. 90% of xDSL consumers experience only 1% or less packet loss. It should also be noted that xDSL has a slightly wider distribution than cable and FTTx, whose distributions are virtually identical. This is a strong indication that the vast majority of consumers experience very low packet loss.

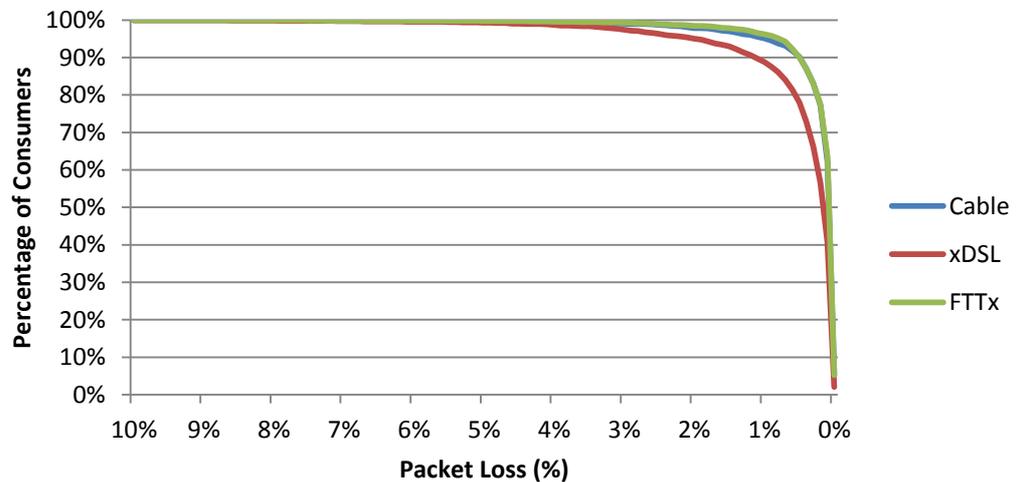


Figure EU.1-18: Cumulative Distribution of Packet Loss, by technology

c.1.6 **DNS Resolution and Failure Rate**

DNS Resolution Time

Figures EU.1-19 and EU.1-20 shows DNS resolution time during the peak and 24-hour periods. As with latency and packet loss, DNS resolution time is much higher for xDSL based services compared to cable and FTTx. This is to be expected as DNS is directly affected by the round-trip latency of the underlying technology. DNS response times are slightly higher at peak hours, Cable and FTTx technology exhibit a much smaller change during peak hours than in October 2013, resulting in a drop in response times during peak hours. This is not the case for xDSL technology, which exhibits higher DNS response times during both periods.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
DNS Resolution (ms)								
Oct-14	37.62	36.3	17.95	17.03	17.81	17.52	25.89	25.01
Oct-13	37.07	35.92	18	16.91	18.35	18.15	27.02	26.15

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

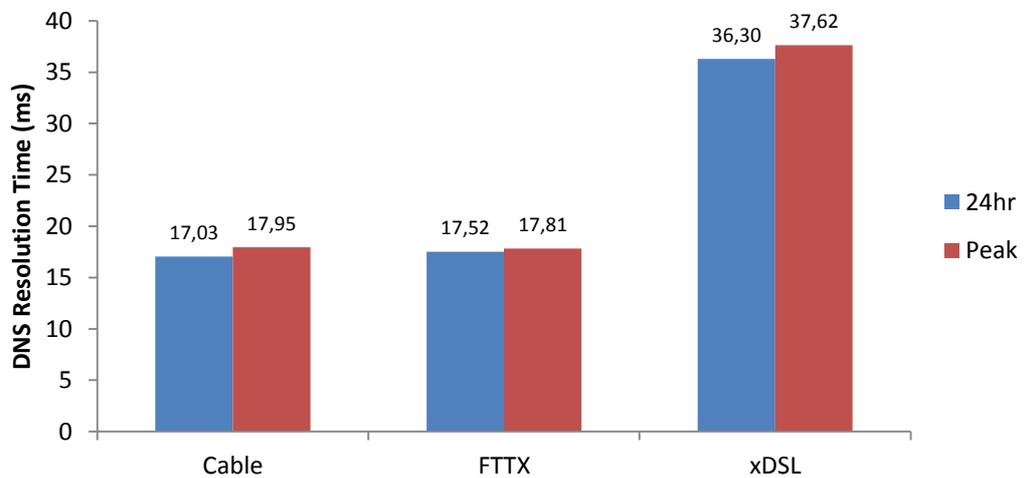


Figure EU.1-20: Peak Period and 24-hour DNS Resolution Time, split by technology (lower is better)

Figure EU.1-21 below displays DNS resolution time split by time of day and access technology. As indicated in figure EU.1-20, FTTx remains very consistent

throughout the day, experiencing almost no change during peak hours. DNS resolution times for cable are lower than for FTTx earlier in the day, increasing during the afternoon period to match FTTx. xDSL also increases steadily throughout the day and more sharply during peak hours.

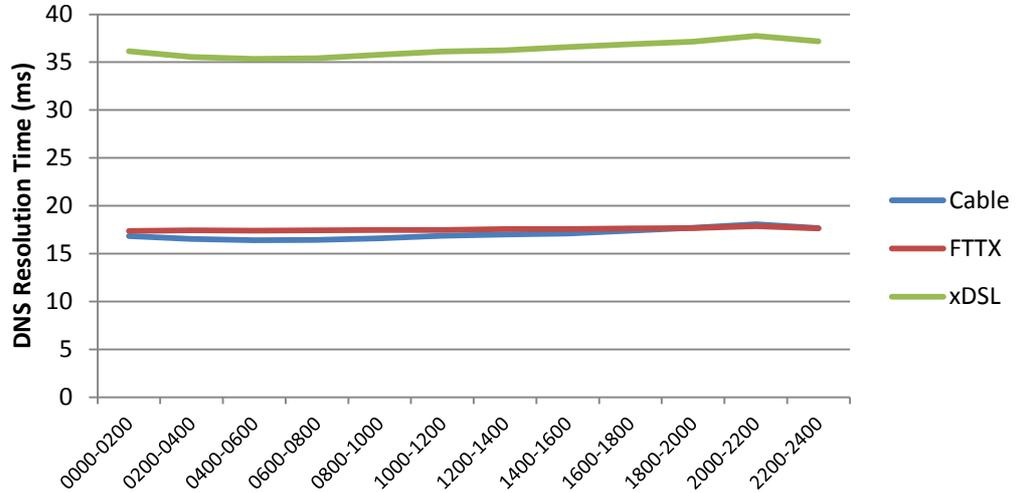


Figure EU.1-21: DNS Resolution Time, split by hour of day and technology (lower is better)

Figure EU.1-22 shows the cumulative distribution plot for DNS resolution split by access technology. The figure below shows virtually no change in DNS resolution from the previous testing period, with 50% of cable and FTTx consumers seeing DNS resolution times of 15ms or less, as was the case in October 2013. DNS resolution is much higher for xDSL users, as was shown above in figures EU.1-20 and EU.1-21 as well.

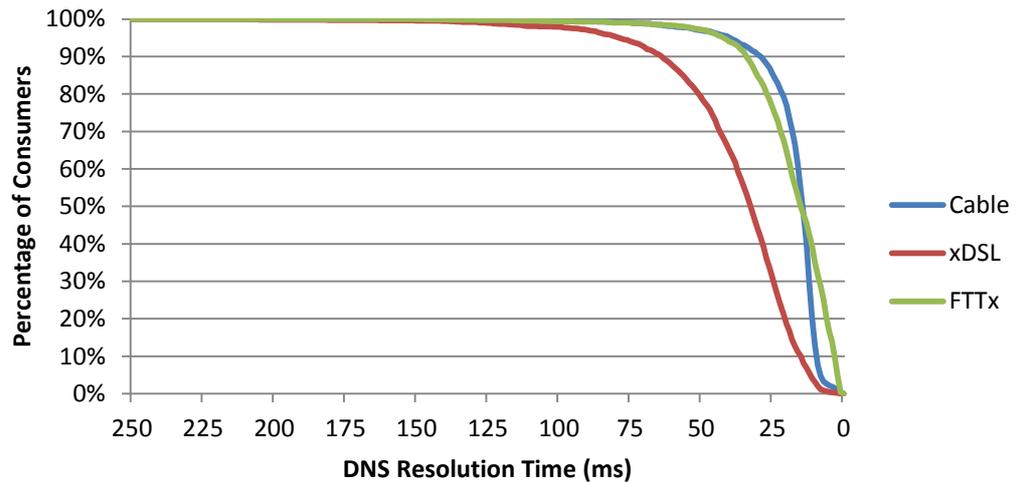


Figure EU.1-22: Cumulative Distribution of DNS Resolution Time, by technology

DNS Resolution Failure Rate

Figures EU.1-23 and EU.1-24 show peak period and 24-hour DNS resolution failure rate across all access technologies. Cable technology displays significant improvements in failure rates since October 2013, showing much lower figures during the peak and 24-hour measurement periods as well as a much smaller change during peak hours, with failure rates proving to be slightly higher during the latter period for all access technologies. FTTx technology also appears to have become more stable, coinciding with the behaviour of DNS response times, although failure rates for this technology are slightly higher than in October 2013. xDSL technology also experiences a slight improvement across both measuring periods.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
DNS Failure Rate (%)								
Oct-14	0.38%	0.30%	0.27%	0.24%	0.24%	0.23%	0.30%	0.26%
Oct-13	0.46%	0.31%	0.56%	0.41%	0.24%	0.21%	0.42%	0.31%

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

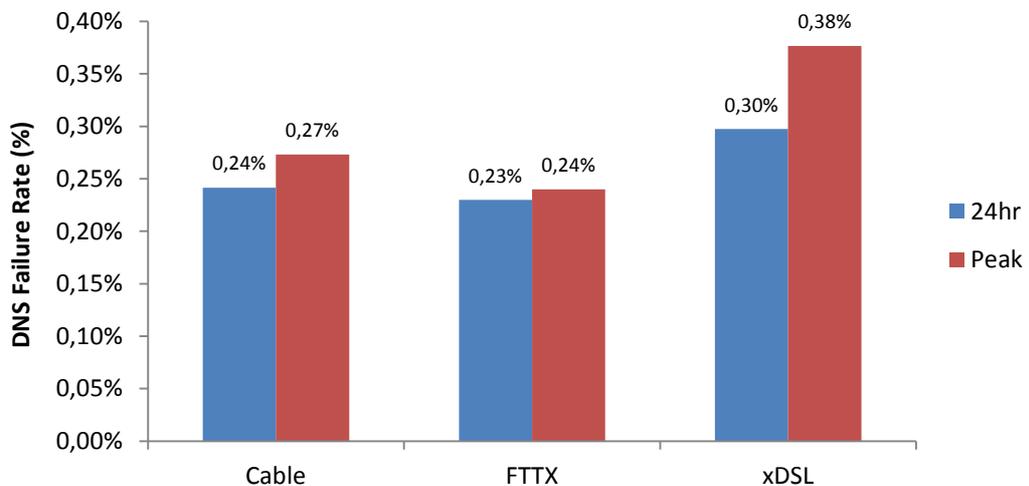


Figure EU.1-24: Peak period and 24-hour DNS Resolution Failure Rate, by technology (lower is better)

Figure EU.1-25 shows DNS resolution failure rate by time of day and split by access technology. This figure shows that the behaviour of failure rates of cable & FTTx technologies is very uneven for all technologies, particularly cable during the early morning period. All access technologies display slightly higher failure rates during the afternoon and peak periods, especially xDSL technology.

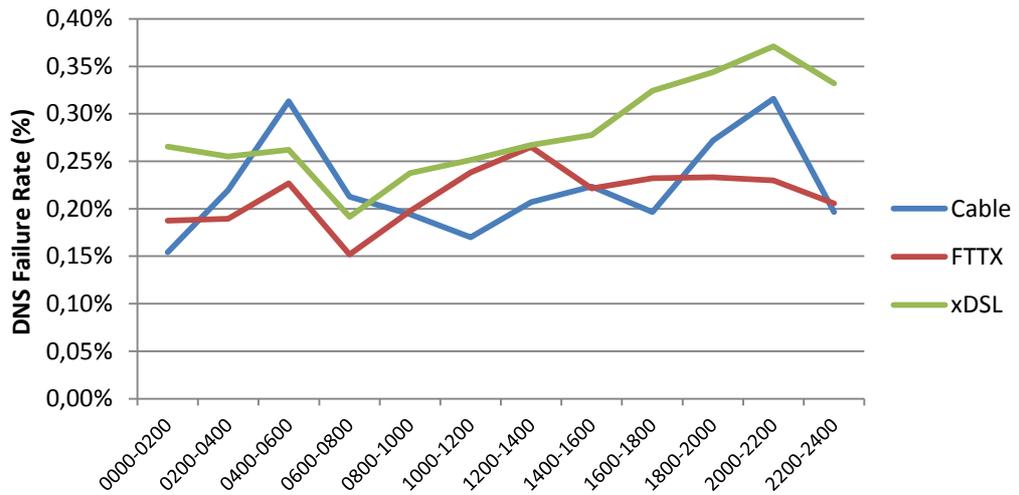


Figure EU.1-25: DNS Resolution Failure Rate, by hour of day and technology (lower is better)

Figure EU.1-26 shows the cumulative distribution chart of DNS resolution failure rate from each type of access technology. As was the case with packet loss, the chart below shows that high failure rates are uncommon among the majority of users for each type of technology, with approximately 90% of users of all access technologies seeing less than 1% failure rates.

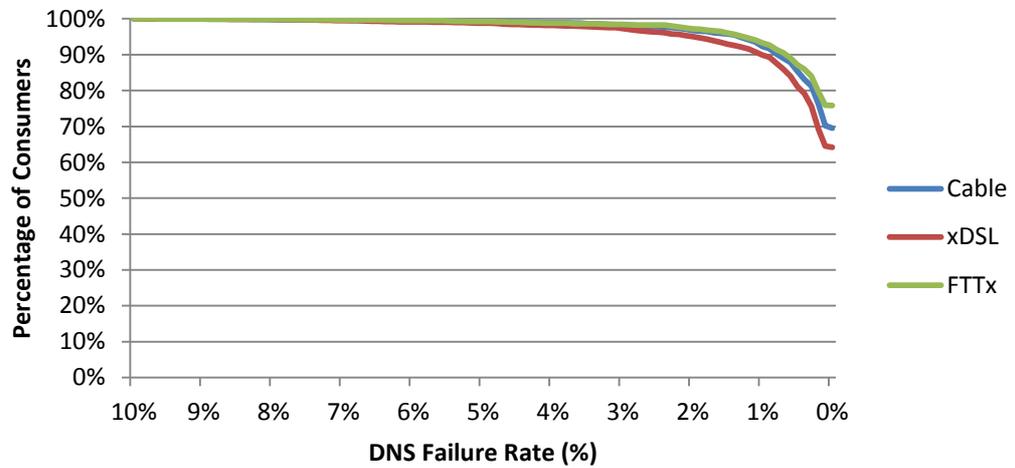


Figure EU.1-26: DNS Resolution Failure Rate, by hour of day and technology (lower is better)

c.1.7 Web Browsing

Figures EU.1-27 and EU.1-28 show the average webpage loading times across all the access technologies covered in this study during the peak and 24-hour periods. It should be noted that this test measures against real websites e.g. 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 exhibits by far the highest webpage loading times during both the peak and 24-hour periods. Web browsing speed for xDSL technology proved to be more than twice as slow as cable and FTTx services, which performed nearly identically during the 24-hour and peak periods, with FTTx proving to be slightly quicker. FTTx services also experienced the smallest increase in loading times during the peak period. Loading times for all access technologies is higher during this report's measurement period compared to October 2013, which may be due to larger page sizes and additional objects on each page.

The behaviour of webpage loading times also mimics the patterns exhibited by latency and DNS resolution time, which is to be expected given that web browsing performance is a function of both line speed and round trip latency. Additionally, beyond a certain level of downstream throughput, usually 10Mbps, latency becomes the main element affecting web browsing performance. This is shown by results in the figures below.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Webpage Loading Time (s)								
Oct-14	2.21	2.06	0.85	0.79	0.83	0.78	1.4	1.3
Oct-13	1.7	1.64	0.64	0.6	0.62	0.6	1.13	1.08

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

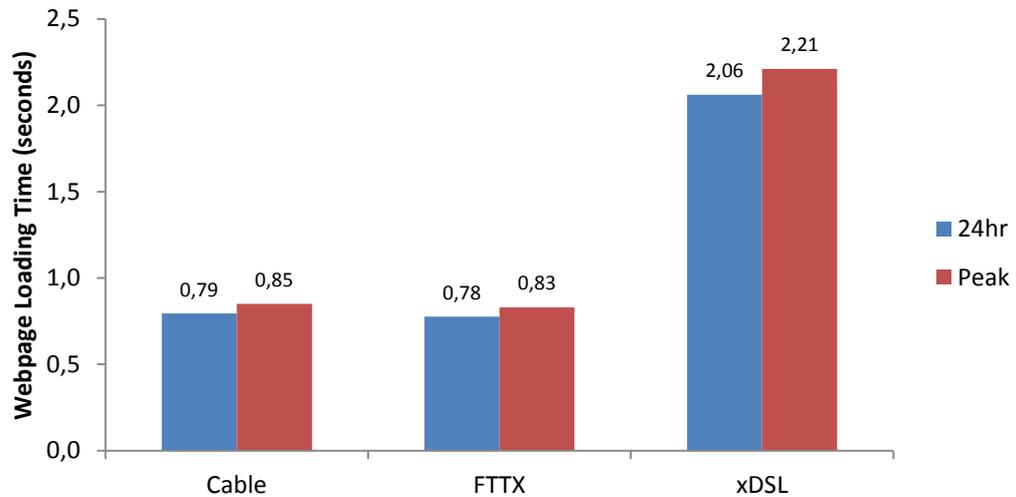


Figure EU.1-28: Peak period and 24-hour Webpage Loading Time, by technology (lower is better)

Figure EU.1-29 below shows the average webpage loading time by time of day and access technology. All technologies exhibit very similar patterns to latency throughout the day. As was the case in October 2013, FTTx technology displays the most consistent loading times, exhibiting the smallest increase during peak hours compared to other access technologies. Peak hour changes are more evident for cable and xDSL technologies, particularly the latter which also shows much higher webpage loading times.

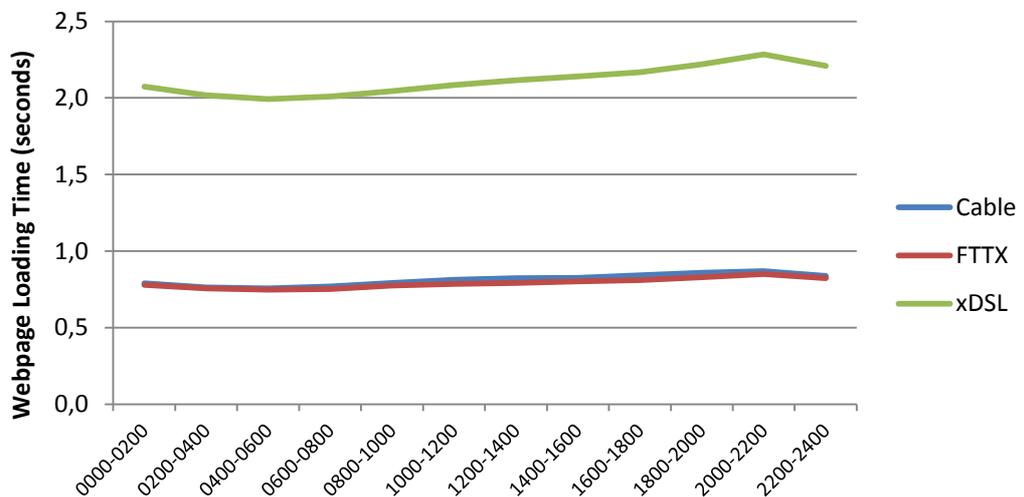


Figure EU.1-29: Webpage Loading Time, by hour of day and technology (lower is better)

Figure EU.1-30 displays the cumulative distribution plot for webpage loading time split by technology. 90% of cable and FTTx consumers experience loading times of

1.2 seconds at most, slightly higher than in October 2013, with both types of technologies exhibiting virtually identical distribution across their respective user bases. 90% of users of xDSL technology, which has a wider distribution than cable and FTTx technologies, see loading times of approximately 3.8 seconds or less, also higher than in the previous year.

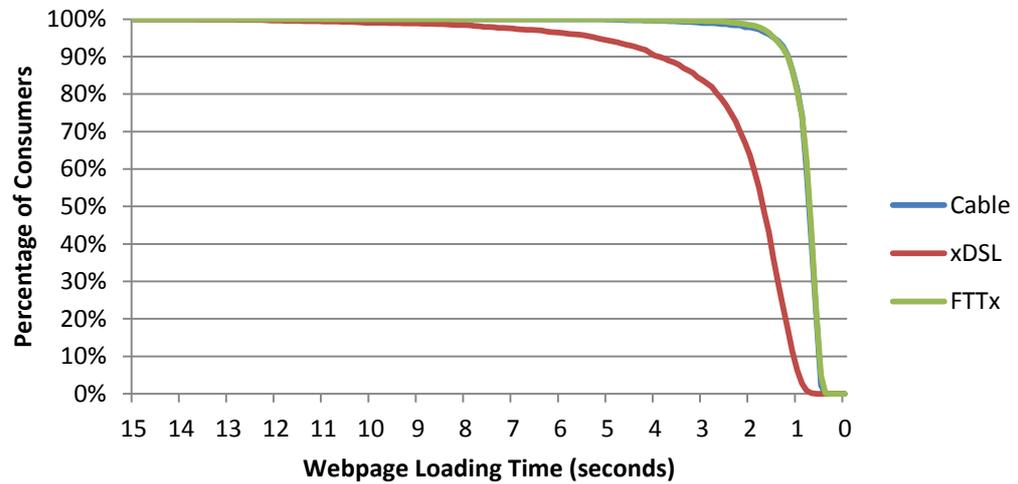


Figure EU.1-30: Cumulative Distribution of Webpage Loading Time, by technology

C.1.8 **VOIP**

Downstream VoIP Jitter

Figures EU.1-31 and EU.1-32 show downstream jitter during the peak and 24-hour periods across all types of access technology. Downstream jitter has improved slightly for cable technology since October 2013, displaying lower jitter during the 24-hour period. All other technologies exhibit a slightly higher level of jitter, particularly xDSL technology which also experiences the largest increase in downstream jitter during peak hours.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Downstream Jitter (ms)								
Oct-14	1.11	0.79	0.59	0.39	0.7	0.46	0.84	0.58
Oct-13	0.96	0.73	0.59	0.41	0.59	0.42	0.76	0.56

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

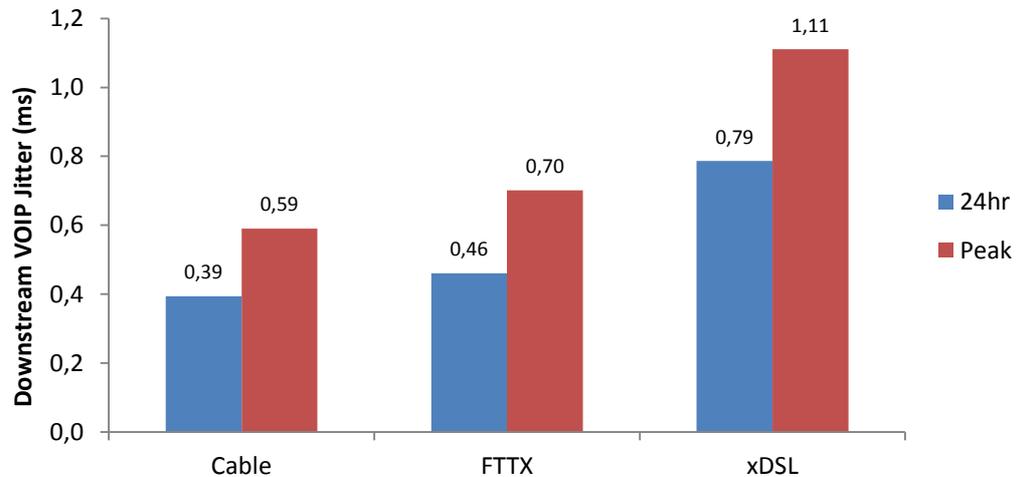


Figure EU.1-32: Peak period and 24-hour Downstream VoIP Jitter, by technology (lower is better)

Figure EU.1-33 shows downstream jitter split by hour of day and access technology. The behaviour of jitter for all technologies is very similar, increasing steadily throughout the day followed by sharper increases during the peak period. Cable technology displays the lowest level of downstream jitter compared to other access technologies, contrasting with downstream jitter for xDSL services, which is noticeably higher and increases more sharply.

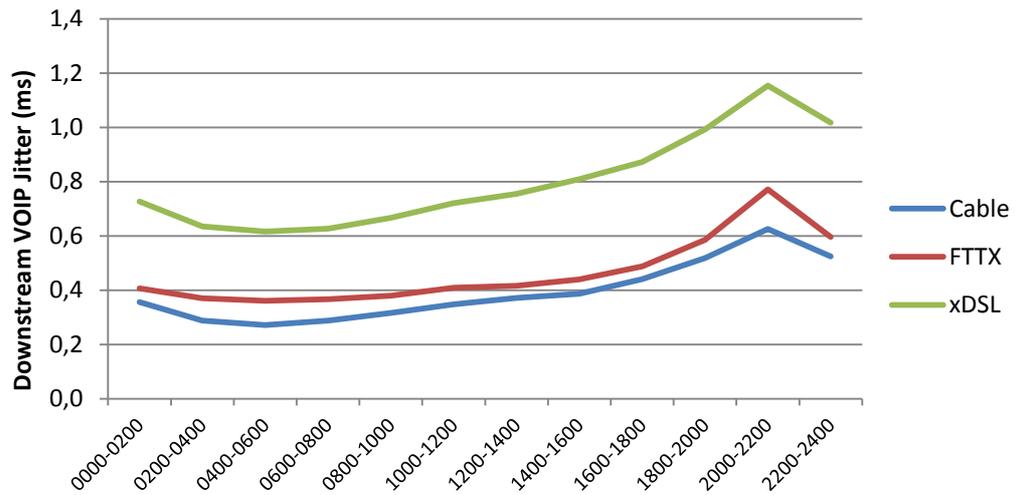


Figure EU.1-33: Downstream VoIP Jitter, by hour of day and technology (lower is better)

The cumulative distribution for downstream VoIP jitter, shown below in figure EU.1-34, shows a tight distribution across all access technologies covered in this study, particularly cable technology, as was the case for web browsing speed. In tandem with average performance, the distribution for xDSL and FTTx technology also appears to have widened slightly since October 2013. 80% of cable and FTTx consumers experience downstream jitter of 0.7ms and 1ms respectively or less, with 80% of xDSL users seeing jitter of 1.5ms or less.

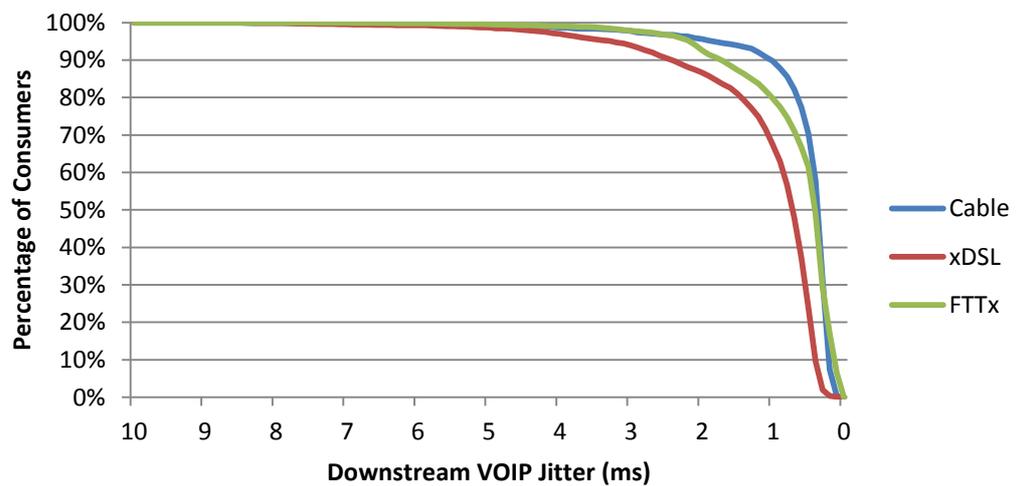


Figure EU.1-34: Cumulative Distribution of Downstream VoIP Jitter, by technology

Upstream VoIP Jitter

Figures EU.1-35 and EU.1-36 show upstream VoIP jitter for the access technologies covered in this study during the peak and 24-hour periods. As was the case in October 2013, cable products exhibit the highest level of upstream jitter across all access technologies and the greatest increase during peak hours. FTTx services never exceeded 1ms of upstream jitter, and also experienced the smallest amount of increase during the peak period. Although the upstream jitter of cable technology is much higher than that of all other access technologies, it is also the only one to display a slight improvement in average performance. FTTx and xDSL technologies instead exhibit a slightly higher level of upstream jitter since the previous year.

The reason for cable services exhibiting higher upstream jitter 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. If the modem is busy whilst the user tries to send a packet, 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 dejitter buffer of at least 25ms, meaning jitter under 25ms would not affect 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	EU Peak	EU 24hr
Upstream Jitter (ms)								
Oct-14	2.1	1.83	3.16	2.51	0.87	0.76	1.96	1.64
Oct-13	1.77	1.58	3.37	2.69	0.86	0.75	1.92	1.62

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

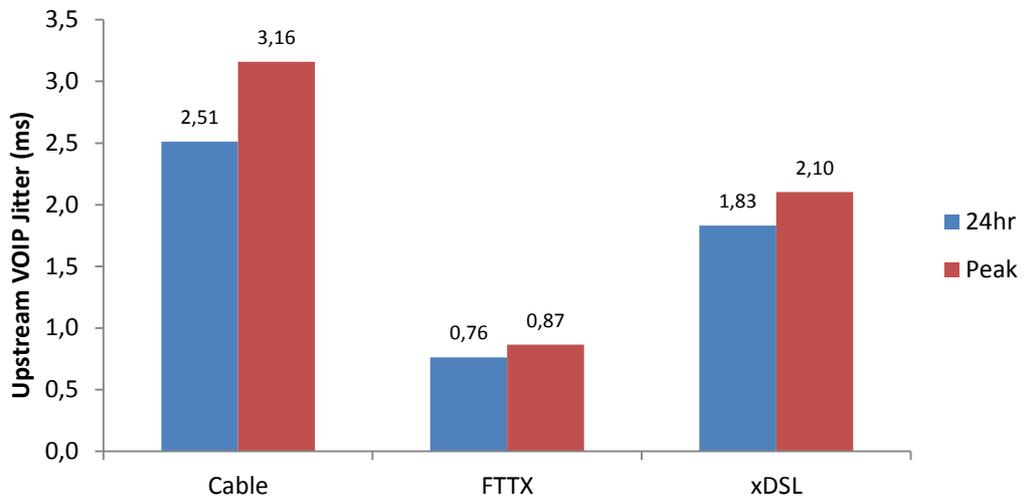


Figure EU.1-36: Peak period and 24-hour Upstream VoIP Jitter, by technology (lower is better)

Figure EU.1-37, which depicts upstream jitter by hour of day and technology, shows that the behaviour of jitter across all access technologies does not differ significantly from the previous measurement period. Upstream jitter for FTTX and xDSL remains mostly consistent, slightly increasing throughout the day and increasing more sharply during the peak period. This is more evident in xDSL technology. Cable services exhibit a much more severe increase in jitter.

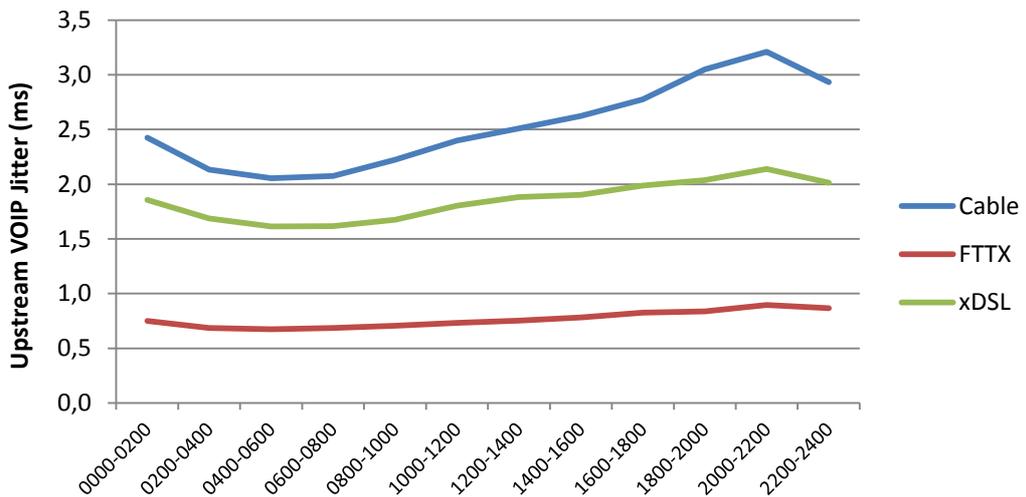


Figure EU.1-37: Upstream VoIP Jitter, split by hour of day and technology (lower is better)

Figure EU.1-38 shows the cumulative distribution of upstream jitter for all access technologies. The distribution of figures is again very similar to what was observed in October 2013. All technologies, particularly FTTx and xDSL, have a tight distribution. Cable services also display a tighter distribution compared to the last measurement period. 80% of cable users experience 3.6ms or less of upstream jitter, as opposed to 4ms in the previous year. 80% FTTx and xDSL see approximately 1ms and 2ms of jitter respectively, as was the case in October 2013.

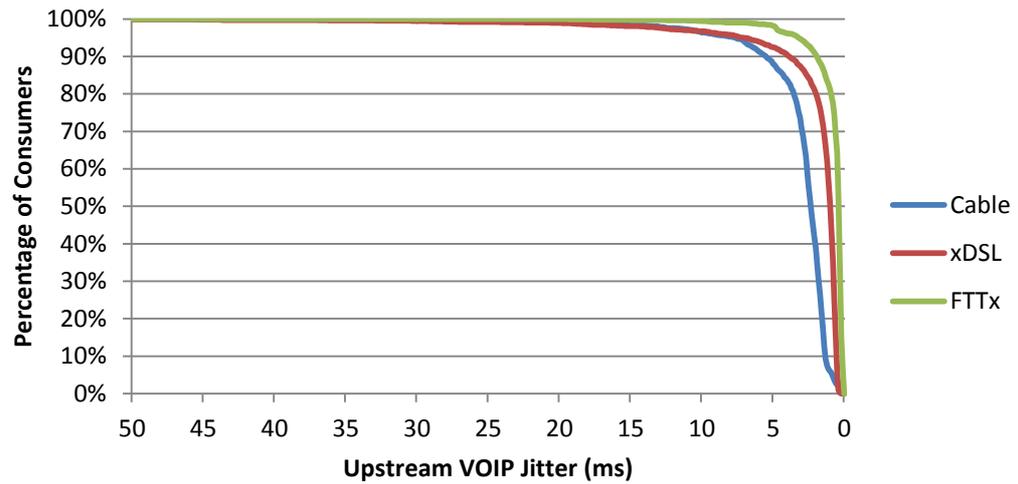


Figure EU.1-38: Cumulative Distribution of Upstream VoIP Jitter, by technology

c.2 Comparison with the United States³

c.2.1 Download Speed

Figure EU.1-39 below shows the average actual and advertised download speeds in Europe and the US for all access technologies. As was the case in October 2013, while throughput expressed as a percentage of advertised speed is lower, actual download throughput in Europe is typically better than in the USA across all access technologies, especially in the case of cable. This suggests that significant differences in the way products are advertised to the general public exist between the two regions.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual/Advertised (%)	US Advertised (Mbps)	US Actual (Mbps)	US Actual/Advertised (%)
xDSL	14.04	8.27	63.32%	8.41	7.67	91.19%
FTTx	64.92	53.09	83.14%	36.53	41.35	113.19%
Cable	79.80	66.57	86.51%	24.88	25.48	102.42%

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

c.2.2 Upload Speed

Figure EU.1-40 shows the actual and advertised upload throughput for all access technologies in Europe and the USA during the peak period. The performance of FTTx and xDSL technologies between the two regions is not significantly different. In contrast, cable technology performance is much better in Europe. As was the case with download speed, the USA shows much better throughput expressed as a percentage of advertised speed, with cable and FTTx technology exceeding the advertised speed. This again indicates differences in marketing schemes between the two regions.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual/Advertised (%)	US Advertised (Mbps)	US Actual (Mbps)	US Actual/Advertised (%)
xDSL	1.05	0.85	82.70%	0.88	0.86	98.2%
FTTx	27.91	25.23	92.29%	22.79	25.93	113.79%
Cable	8.43	8.33	100.10%	4.07	4.51	110.71%

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

³ Data taken from Measuring Broadband America – April 2014 - <http://www.fcc.gov/measuring-broadband-america>

c.2.3 Latency

Figure EU.1-41 shows a comparison of latency figures between Europe and the USA, split by access technology. Contrasting with results from the previous measurement period, Europe displays much better latency across all access technologies. There is a significant difference in latency for both cable and xDSL, but not for FTTx. This again suggests xDSL and cable based services are deployed in a similar way throughout each region.

Technology	Europe	US
xDSL (ms)	37.36	48.96
FTTx (ms)	20.16	24.07
Cable (ms)	19.22	31.77

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

c.2.4 Packet Loss

Figure EU.1-42 is the comparison of packet loss during the peak period between Europe and the USA, split by technology. As was the case in October 2013, all access technologies in the USA displayed significantly lower packet loss compared to Europe. In actuality, the difference is not significant and is negligible with respect to broadband performance for individual users.

Technology	Europe	US
xDSL	0.45%	0.28%
FTTx	0.21%	0.18%
Cable	0.23%	0.12%

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

D Comparison Between Countries

D.1 Key Performance Indicators⁴

D.2 Download and Upload Speeds

D.2.1 Download

The metric most commonly associated with broadband performance is download throughput, and it is also the metric that ISPs typically use to advertise their products. Because of this, it receives a large amount of attention from regulators and ISPs when marketing their products.

In order to compare between countries and technologies, which can often have very different performance characteristics, results for download speed are presented as a percentage of advertised speed. This is done primarily to allow the reader to determine more easily how accurate marketing claims of ISPs in certain countries are.

Figures EU.2-1 to EU.2-3 below represent download speed as a percentage of advertised speed for each country considered in this study for xDSL, FTTx and cable technologies. Each figure also shows the average achieved across all countries included in this study. For all technologies, it is typical for most countries to achieve a level of throughput close to the average with few outliers, particularly with regards to xDSL technology. Cable and FTTx perform very well across all countries that meet a minimum sample size, achieving 86.18% and 84.05% of advertised speed on average respectively. Almost all countries using either technology generally perform close to the overall average. Malta is seen to perform the furthest below the average compared to all other countries, although it still remains relatively close to it with 74.48%. This contrasts with the previous measurement period when Malta exceeded its advertised speed, even though there have been no changes to the test conditions.

xDSL technology exhibits a wider distribution of results. Slovakia is again shown to achieve the highest level of throughput as a percentage of advertised speed, with the UK proving to have the lowest performance with 44.96%, far below the average of 70.75%. However, as will be shown later in this study, this has more to do with the advertised rate as opposed to actual performance of broadband in these countries. Most countries' throughput performance as a percentage of advertised speed is closer to the average for FTTx and cable technologies.

⁴ The data has not been weighted, however the data has been trimmed as per section B.1.4.2

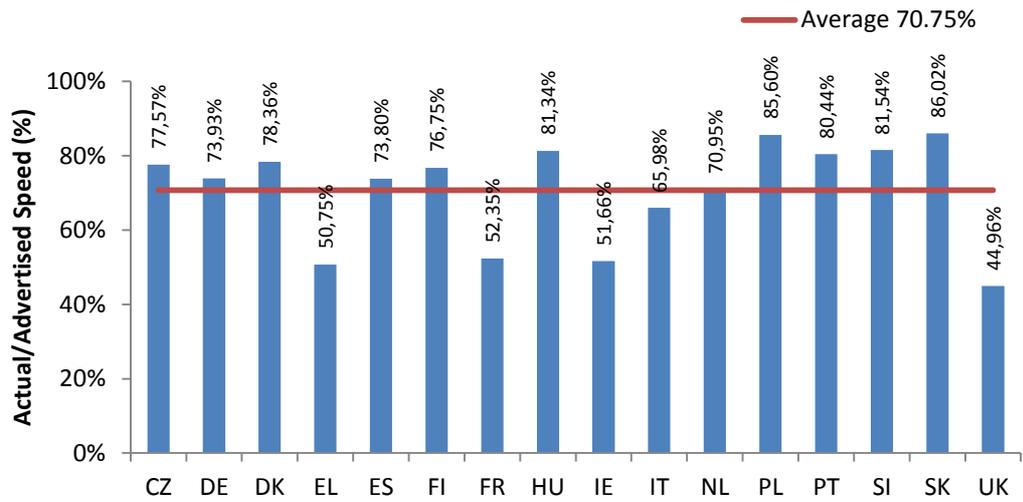


Figure EU.2-1: Actual Download Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

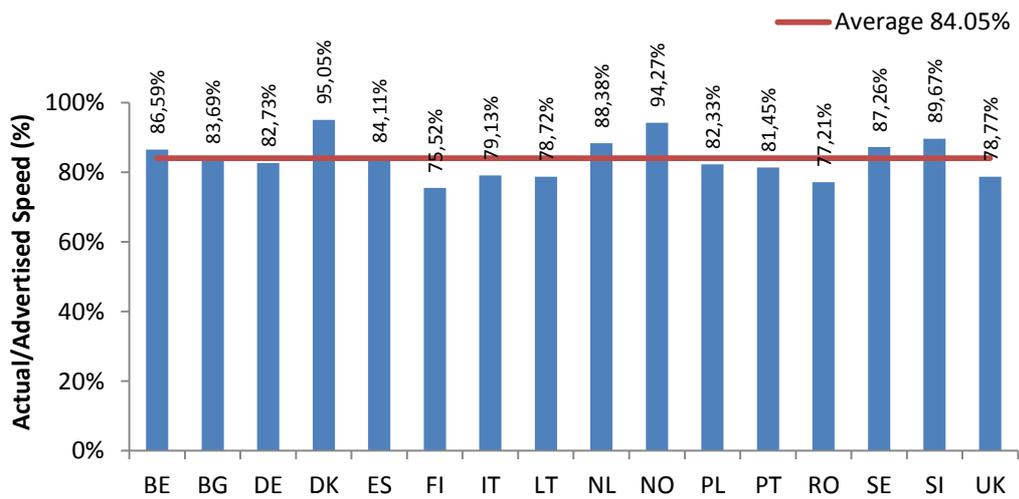


Figure EU.2-2: Actual Download Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

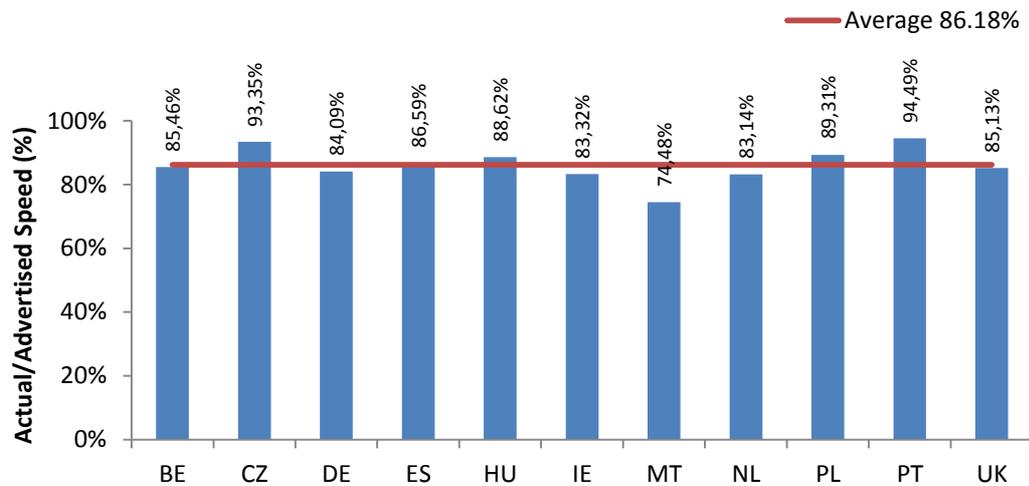


Figure EU.2-3: Actual Download Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

As mentioned above, nations where cable services are more common tend to achieve figures closer to the advertised rate. While xDSL is quite common throughout Europe, it is much more likely to deliver a lower level of performance compared to cable and FTTx. This is due to access speed degrading with increasing copper loop lengths. Countries such as the UK and France, which deliver some of the poorest results for xDSL, perform better for other technologies, with the UK also exceeding the average actual download speed and outperforming all countries apart from Ireland for cable.

Overall averages for all access technologies experience an improvement from the previous measurement period of October 2013, particularly cable which shows a 32.7% increase from 47.84Mbps to 63.49Mbps.

Figures EU.2-4 to EU.2-6 below show actual throughput achieved in each country for all access technologies considered in this study. Only countries and technologies with a statistically representative sample are included. The spread of results for actual speed is significantly wider compared to throughput expressed as a percentage of advertised speed, suggesting that similarities in results seen in figures EU.2-1 to EU.2-3 are due to differing marketing strategies in each country. This is showcased particularly by a comparison of actual throughput between the UK and Slovakia for xDSL, who achieve the lowest and highest level of throughput as a percentage of advertised speed respectively although the UK noticeably outperforms Slovakia in real terms. Slovakia also displays one of the lowest levels of actual throughput across all countries.

In the UK and France, xDSL products are predominantly advertised with a single headline speed. Customers use copper phone lines, meaning they can only receive a fraction of the speed advertised by the package. Other countries will offer a wider array of packages and may adopt policies prohibiting providers from selling products customers cannot achieve full speed on.

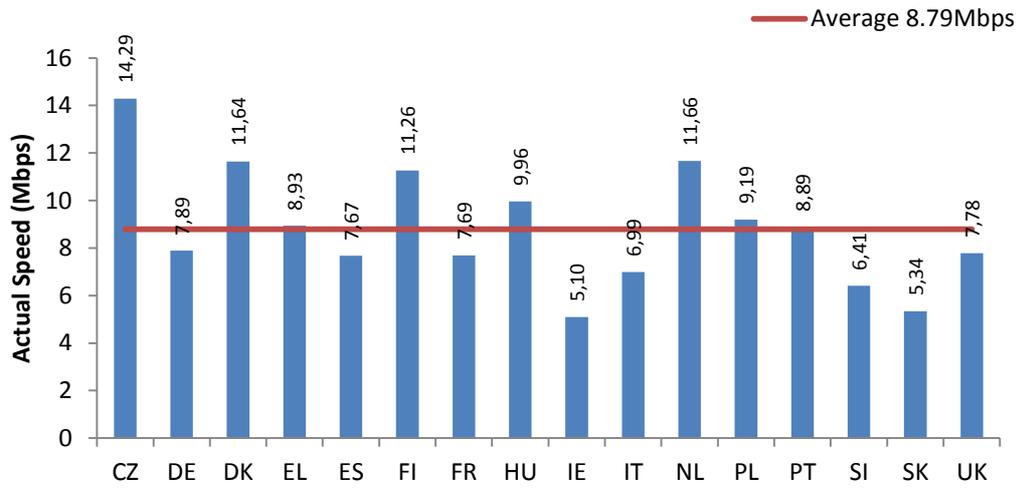


Figure EU.2-4: Actual Download Speed of xDSL technology during peak periods, by country



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

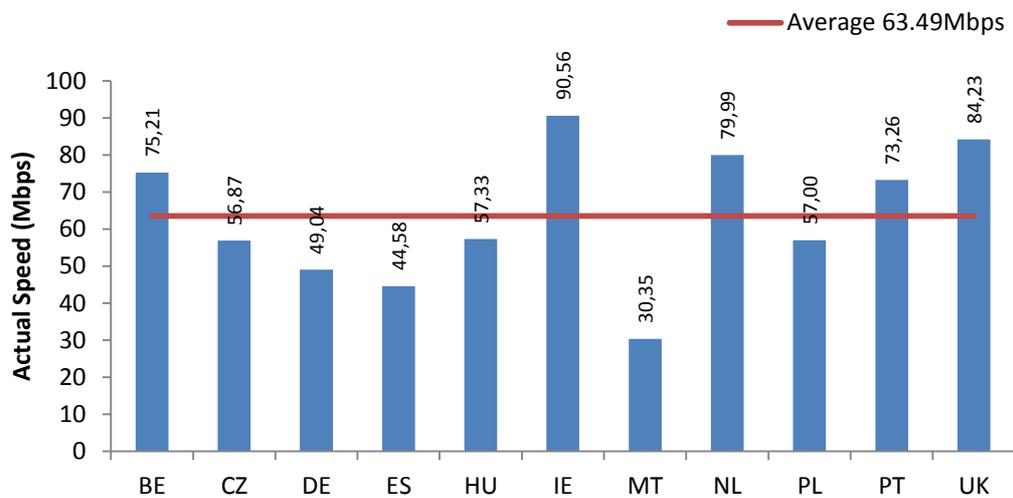


Figure EU.2-6: Actual Download Speed of cable technology during peak periods, by country

D.2.2 Upload

A great majority of European ISPs provide an asymmetric broadband service, with download throughput usually being many times higher than upload throughput. Historically, this has made sense to ISPs who observed users consumed (downloaded) much more content than they produced and shared (uploaded). However, upload speed is gradually becoming a more important metric as more and more users upload photos, videos and use online storage services. Many ISPs today offer higher upload speed services, recognizing this growing trend.

As with download speed, upload results are presented as a percentage of advertised speed first in order to provide a better comparison between different countries.

Figures EU.2-7 to EU.2-9 display upload speed expressed as a percentage of advertised speed for all access technologies. Only countries and technologies with a statistically representative sample are included. The main thing to note is all technologies achieve on average a greater percentage of advertised upload speed, with the average for cable technology performing just under the advertised speed across all nations. This is likely due to the asymmetry of throughput rates (a service needs to handle less traffic in order to deliver a higher percentage in the upstream direction as the rates are lower). This is particularly important for xDSL services, as the lower upstream target is more manageable even on longer copper phone lines. This may also have to do with lower usage of the upstream direction, although there is not enough data to support this theory.

As mentioned above, cable services perform just below the advertised rate on average, achieving 99.92% of advertised speed. FTTx technology also nearly

achieved its advertised rate on average with 96.26%. All access technologies show a slightly lower average than in October 2013, although as with download throughput, this is mostly related to higher advertised rates than actual performance.

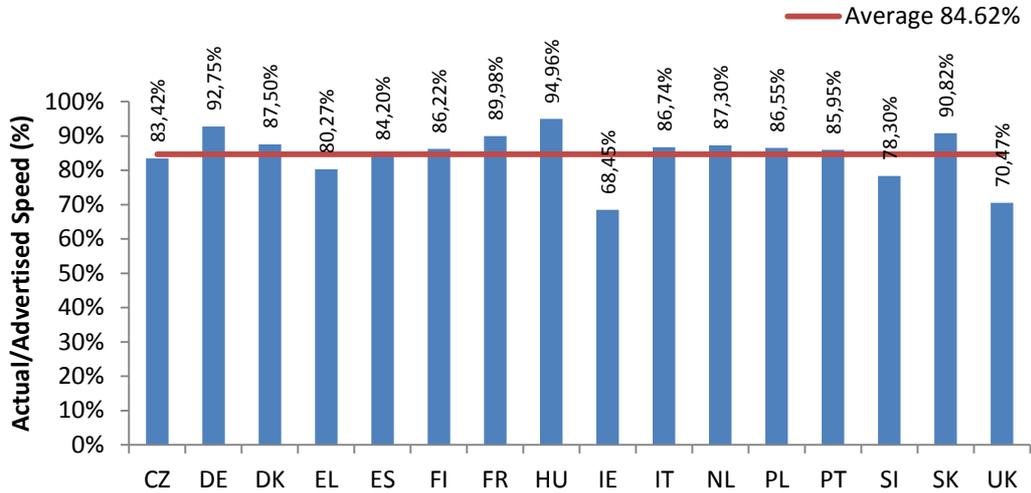


Figure EU.2-7: Actual Upload Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

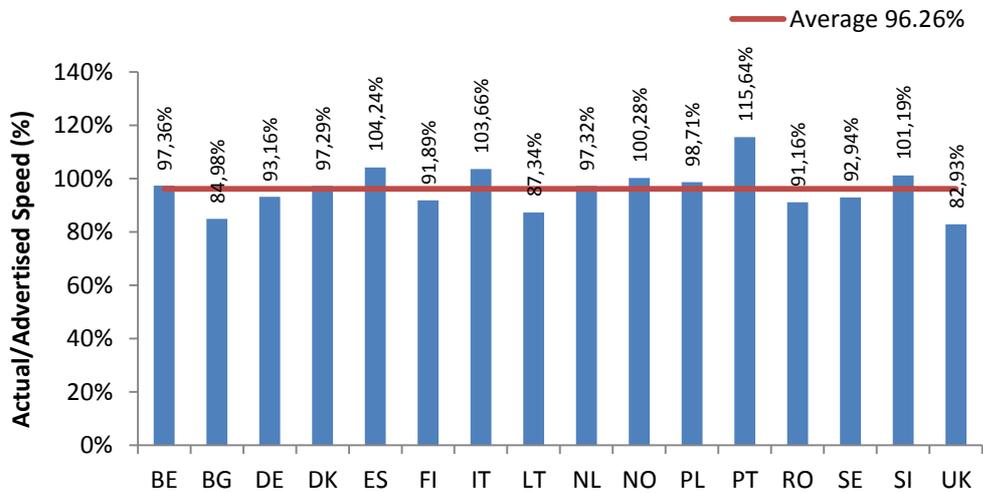


Figure EU.2-8: Actual Upload Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

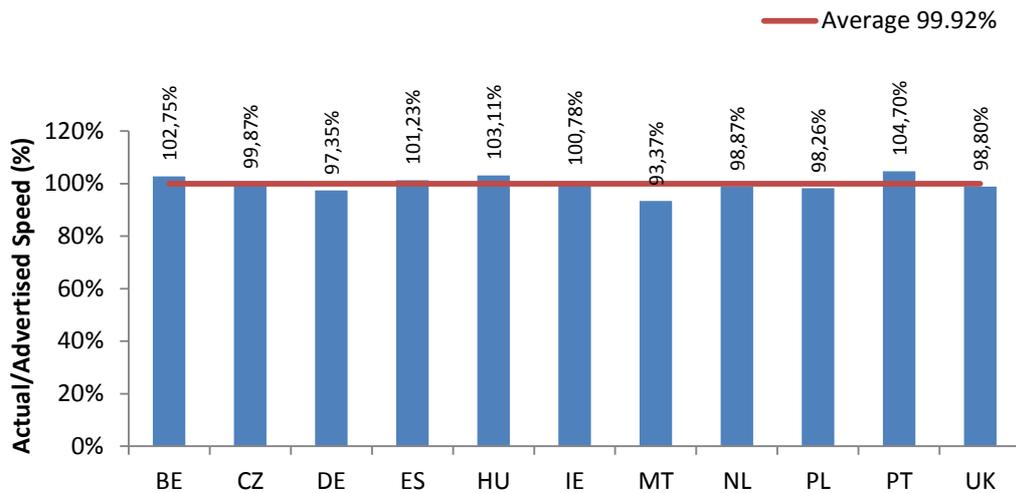


Figure EU.2-9: Actual Upload Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

Numerous countries exceed their individual advertised rates for upload throughput for both cable and FTTx services, with most others performing very close to the advertised rates. Only Bulgaria, Lithuania and the UK reach below 90% of the advertised speed for FTTx technology. All countries using cable technology exceed 90% of the advertised rates. Most countries for xDSL and FTTx technology also exceed the average rate. Portugal especially performed very well across all access technologies, particularly FTTx for which it outperformed all other countries, achieving 115.64% of the advertised speed. It also performs well above the average rate of xDSL and cable technologies.

As with download speed, observing actual results for upload speeds by country and access technology shows an improvement in average performance, contrasting with results for throughput expressed as a percentage of advertised speed.

Figures EU.2-10 to EU.2-12 show actual upload speed for each technology across each country. Eastern European and Nordic countries tend to display higher levels of throughput compared to their western counterparts, as was the case in October 2013.

Lithuania outperforms all other countries significantly with regards to FTTx technology in spite of the fact their performance in upload speed as a percentage of advertised speed was one of the lowest compared to other countries.

As was also observed in October 2013, the asymmetric relationship between download and upload speed is most evident in xDSL technology. With an average download speed of 8.79Mbps and upload speed of 0.86Mbps, the ratio between download and upload is approximately 10:1. Cable technology also shows a ratio of nearly 10:1, contrasting with the previous measurement period, whereas FTTx technology continues to display a ratio of roughly 2:1.

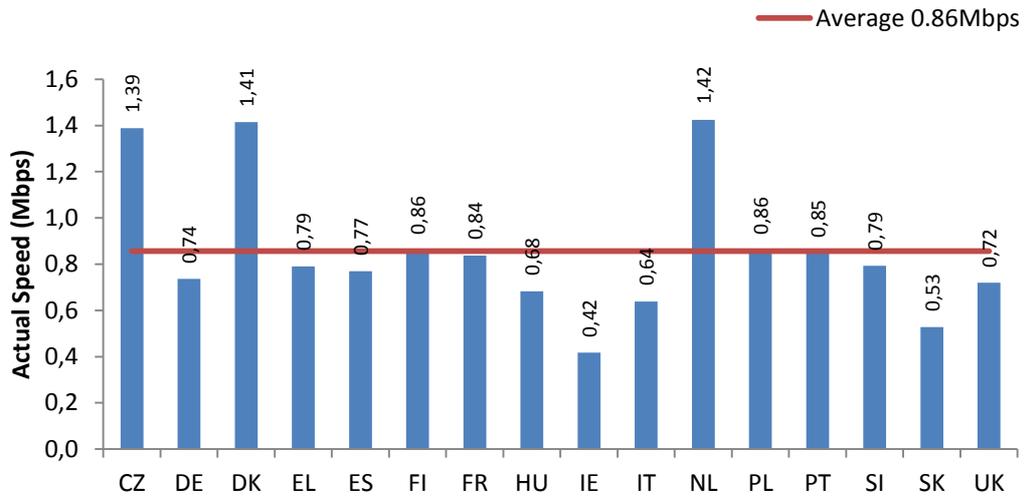


Figure EU.2-10: Actual Upload Speed of xDSL technology during peak periods, by country

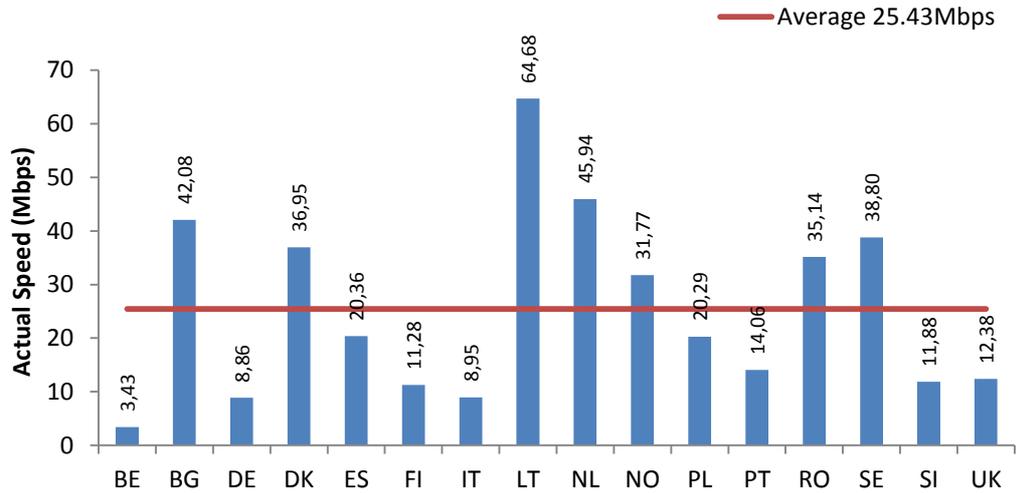


Figure EU.2-11: Actual Upload Speed of FTTx technology during peak periods, by country

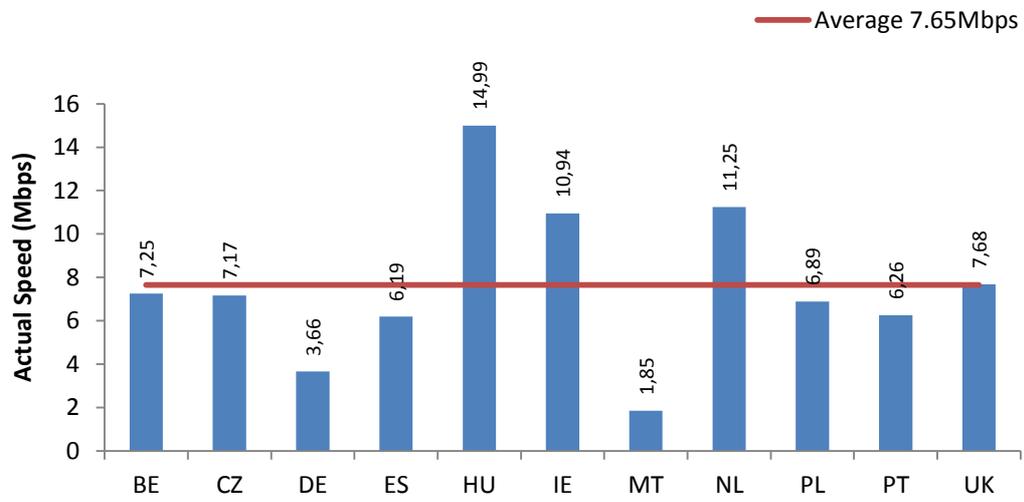


Figure EU.2-12: Actual Upload Speed of cable technology during peak periods, by country

D.2.3 Latency

Latency is an important metric often overlooked by consumers. Latency, also referred to as round-trip latency, is a measure of how long it takes for a single packet of data to go from point A to point B and back. In this study, round-trip latency is measured between the panelists' homes and the nearest measurement server.

Given that every communication with Internet services involves the transmission and reception of packets of data, latency affects everything we do on the Internet. It is especially important for time-sensitive applications such as online gaming, video streaming and voice communications. The lower the latency, the faster and 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 is most typically the dominant factor affecting levels of latency.

Figures EU.2-13 to EU.2-15 show the average round-trip latency per country, split by each access technology. On average, cable delivered the lowest latency of 20.57ms, compared to FTTx's 21.94ms and xDSL's 36.19ms. FTTx technology exhibits a slightly increased average latency since October 2013, with xDSL and cable showing a small improvement.

Thanks to the deployment of FTTH technology in certain eastern European regions, countries such as Bulgaria and Slovenia delivered the best latency performances for FTTx technology. It should also be noted that latency of these two countries also outperforms the latency of all countries using cable technology despite cable achieving a lower overall average. This is because FTTH technology does not need to use an xDSL based last mile technology that causes a significant latency overhead.

Of particular note is the significantly higher latency of xDSL technology from Spain, which also delivered the highest latency results in October 2013. This is despite the fact there are measurement servers located in Spain (Madrid).

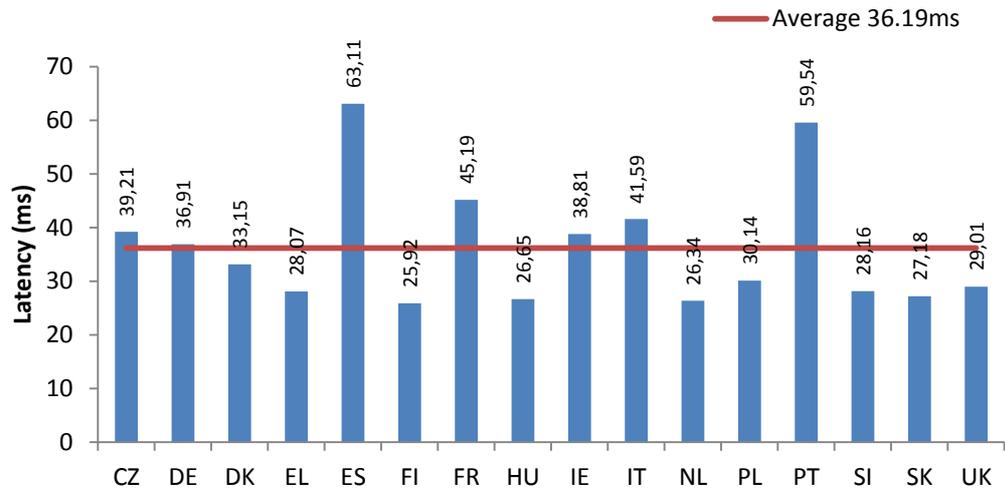


Figure EU.2-13: Latency of xDSL technology during peak periods, by country

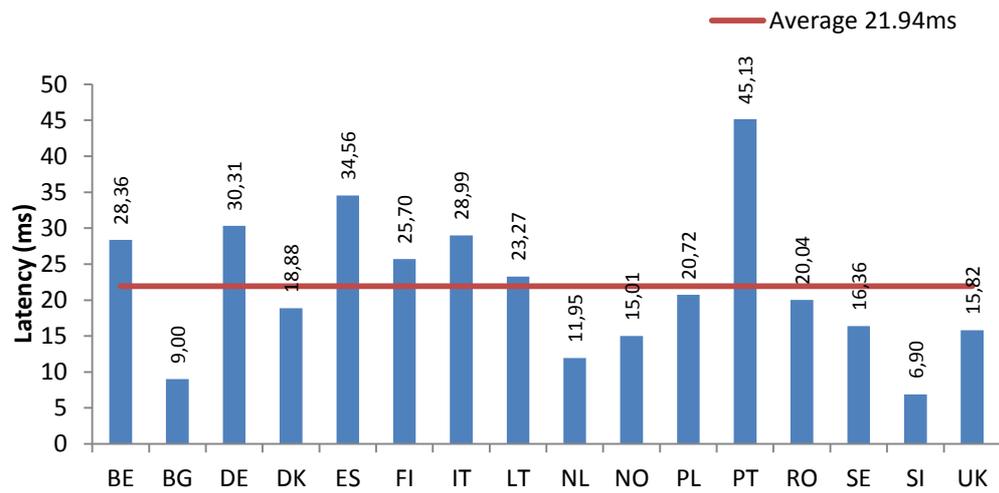


Figure EU.2-14: Latency of FTTx technology during peak periods, by country

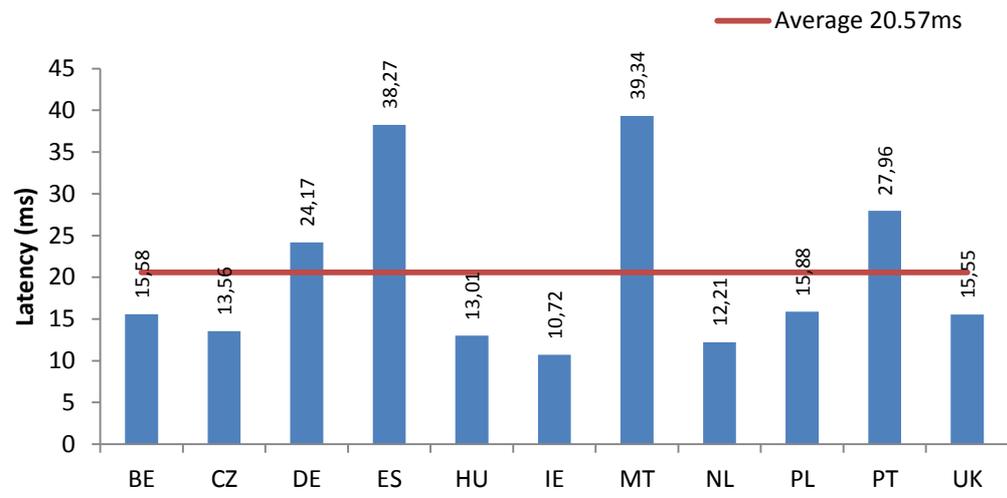


Figure EU.2-15: Latency of cable technology during peak periods, by country

D.2.4 Packet Loss

As can be the case with latency, packet loss is often overlooked by consumers. This metric describes what percentage of packets sent from the home to the server they are communicating with and back are lost during the transmission. When packets are lost, the two parties involved in the communication will usually attempt to retransmit data in order to account for the loss. This takes time and becomes very noticeable to users when it reaches a certain level.

Realtime applications such as online gaming, video streaming and voice communications are the most affected by high packet loss.

Figures EU.2-16 to EU.2-18 show packet loss figures by country, split for each type of access technology during the peak period. As was the case with packet loss during the previous measurement period, most countries exhibit very low packet loss with very few exceeding 0.75%. Most countries shown to exceed 0.5% packet loss are seen on xDSL technology, notably Hungary, Ireland, France and Italy. Italy also shows significantly higher packet loss than all other countries on FTTx technology. FTTx and xDSL technologies both show slightly lower averages since October 2013, in contrast to cable technology which displays a slightly higher average packet loss. This is due to a number of countries exhibiting higher packet loss. Examples include Portugal, Poland and Spain. Differences in packet loss between each country are negligible in real world terms.

It is not unusual for packet loss on xDSL technology to be higher than it is for all other access technologies due to the use of older copper lines. These lines are more likely to suffer from physical defects which may inhibit communications and thus overall performance. Otherwise, most nations across all access technologies show packet loss ranging between 0.1% and 0.4% and also do not exceed the overall averages.

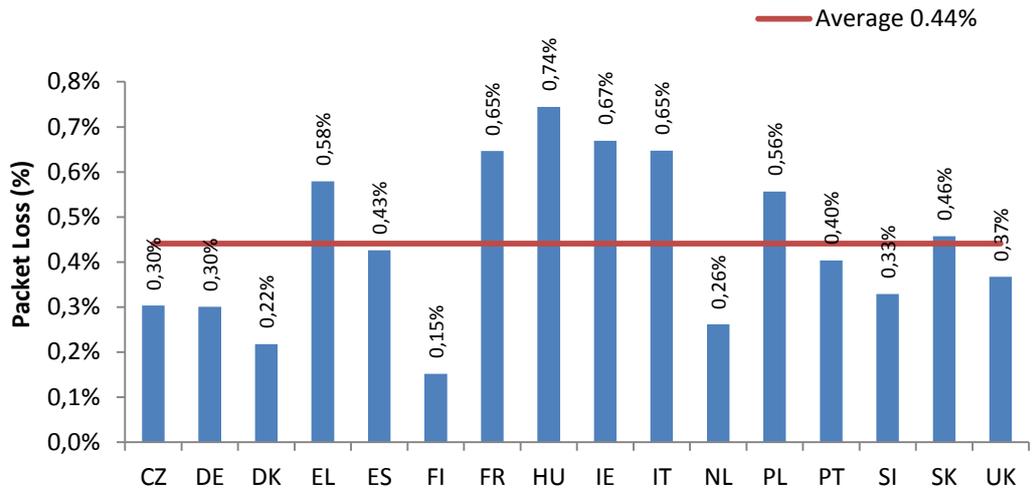


Figure EU.2-16: Packet loss of xDSL technology during peak periods, by country

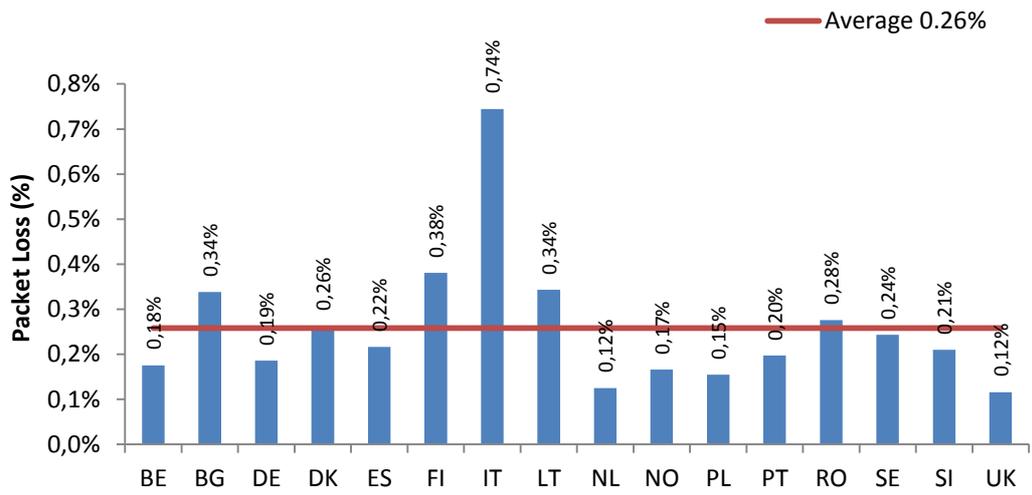


Figure EU.2-17: Packet loss of FTTx technology during peak periods, split by country

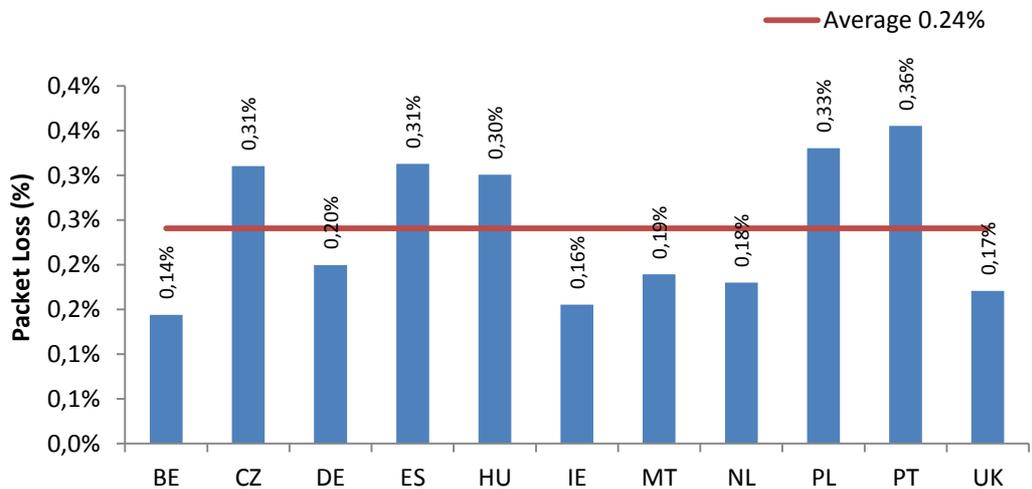


Figure EU.2-18: Packet loss of cable technology during peak periods, by country

D.2.5 DNS Resolution Time and Failure Rate

DNS is a very important Internet service that allows you to turn hostnames, e.g. `www.youtube.com`, into IP addresses that your computer can communicate with. DNS services are typically provided by the ISP to provide a fast, nearby service for their users to use. A DNS service performing badly can lead users to perceive noticeable delays. This is especially evident during web browsing activities, which rely extensively on DNS.

Theoretically, a good DNS deployment should provide DNS resolution times and failure rates that are at most equal to latency and packet loss figures respectively. This is because DNS servers are typically hosted inside the ISP's networks. Thus, this traffic doesn't need to leave the ISP's network.

Figures EU.2-19 to EU.2-21 show DNS resolution time for all technologies split by country during peak hours. As was the case in October 2013, FTTx technology delivered the best DNS resolution times, averaging 17.67ms, with those nations deploying FTTH technology exhibiting the best resolution times of all. Examples include Bulgaria, Lithuania and Slovenia, all of which also outperform countries using cable technology. The average DNS resolution time for cable technology is almost the same at 18.57ms, with xDSL exhibiting a significantly higher average resolution time of 35.06ms.

As was the case in the previous measurement period, Belgium proves to be an exception to other countries using FTTH technology, with its FTTx (VDSL in this case) services averaging much higher DNS resolution times than latency. Other nations whose DNS resolution times don't closely resemble their latency results include Hungary, Italy and Ireland for xDSL technology and Poland for xDSL and FTTx technologies, with DNS resolution times proving to be much higher.

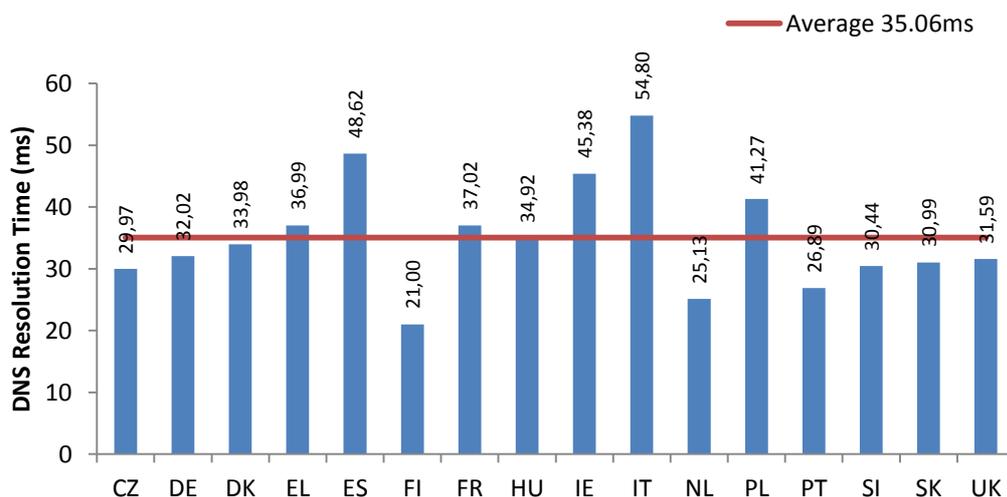


Figure EU.2-19: DNS Resolution Time of xDSL technology during peak periods, by country

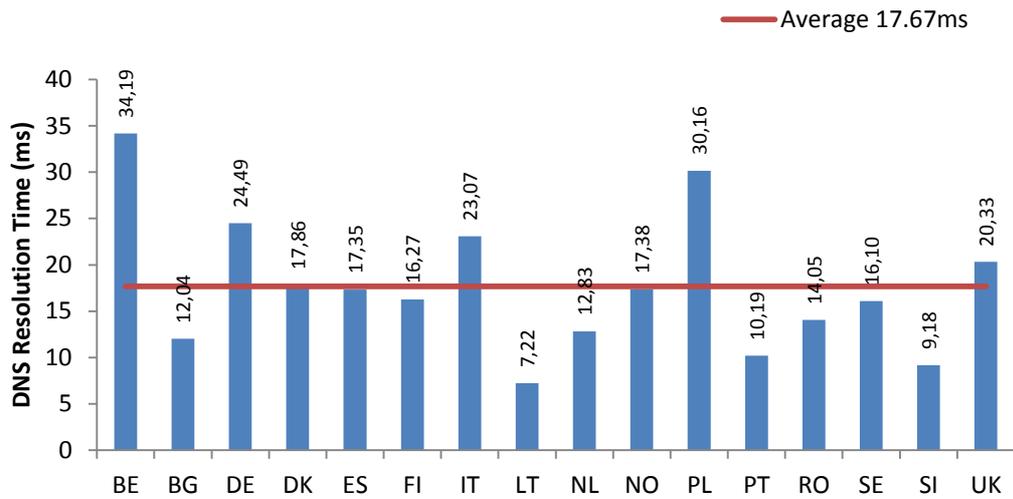


Figure EU.2-20: DNS Resolution Time of FTTx technology during peak periods, by country

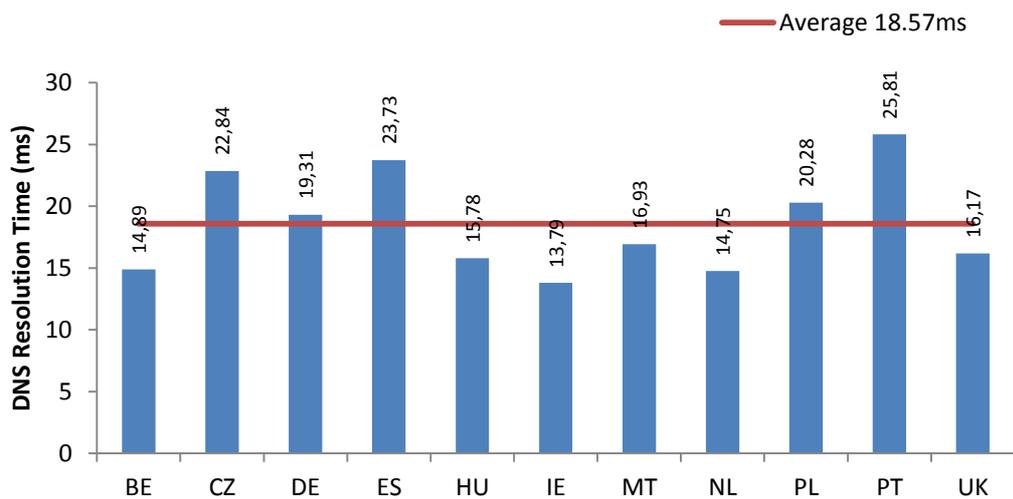


Figure EU.2-21: DNS Resolution Time of cable technology during peak periods, by country

Figures EU.2-22 to EU.2-24 show DNS failure rate per country for each access technology during the peak period. On average, xDSL and cable technologies display lower failure rates than in October 2013, with most countries displaying lower DNS failure rates as well. Exceptions include Slovenia for xDSL technology and Ireland for cable technology.

In contrast, DNS failure rate on FTTx technology is slightly higher on average. Some countries, particularly Bulgaria, Norway and Finland, displayed a significant increase in failure rates, with most other countries reporting a slight increase. Exceptions include Spain and Sweden, which demonstrate a noticeable improvement compared to the previous year.

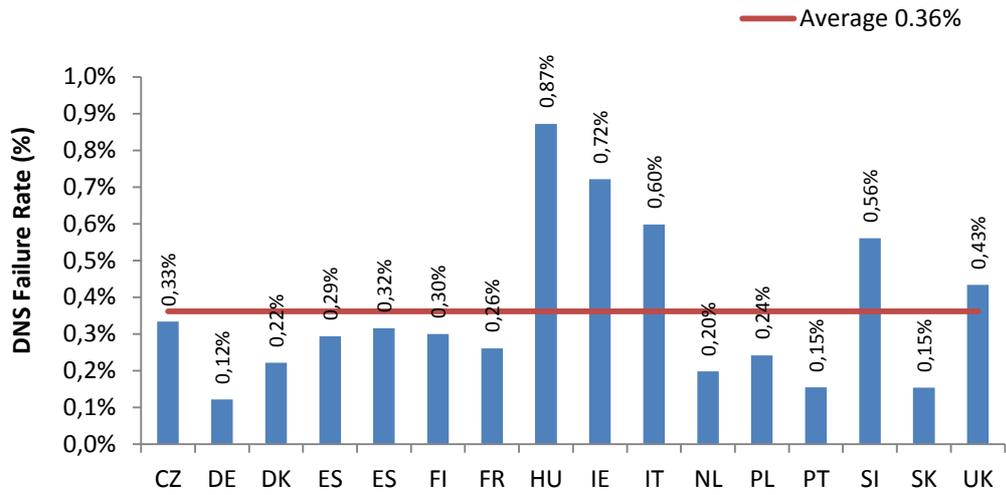


Figure EU.2-22: DNS Resolution Failure Rate of xDSL technology during peak periods, by country

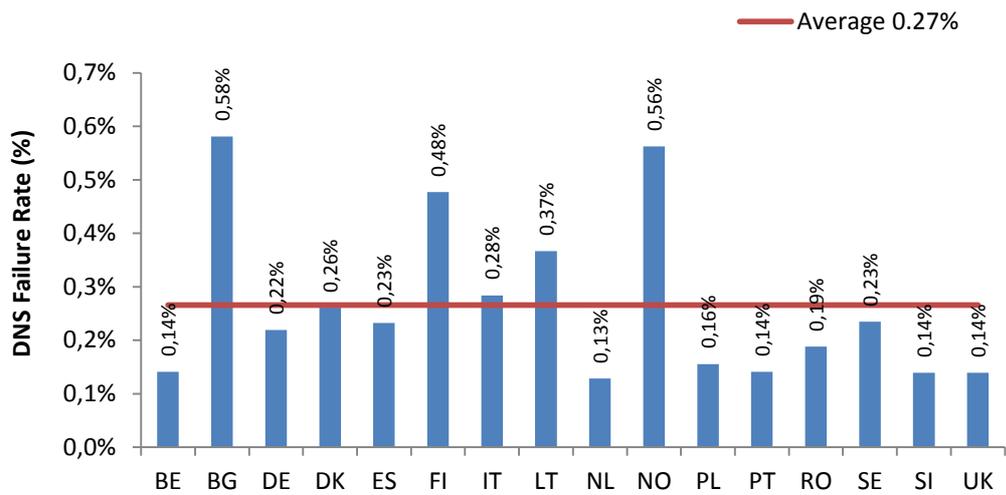


Figure EU.2-23: DNS Resolution Failure Rate of FTTx technology during peak periods, by country

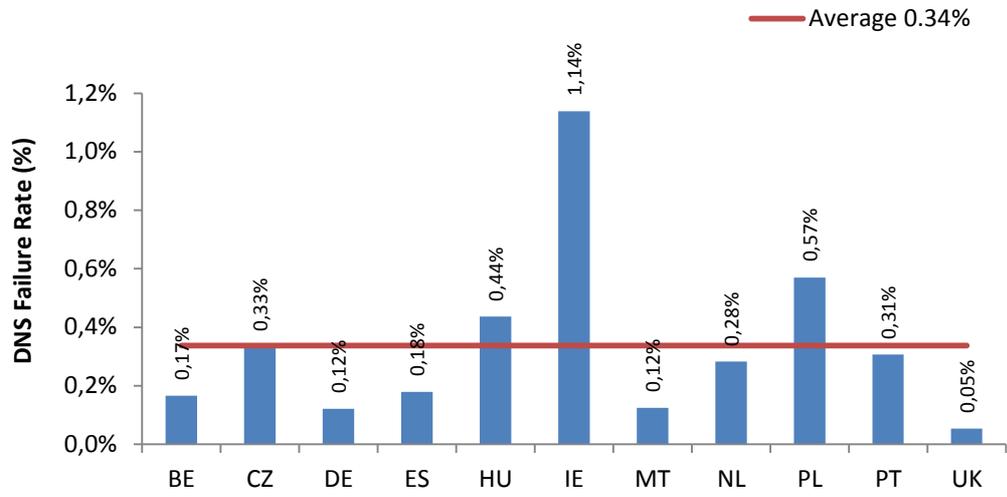


Figure EU.2-24: DNS Resolution Failure Rate of cable technology during peak periods, by country

D.2.6 Web Browsing Speeds

Figures EU.2-25 to EU.2-27 below display webpage loading times in each country during the peak period, split by access technologies. The test was performed to the public-facing websites of Facebook, Google and YouTube, whose servers are geographically hosted across Europe to optimize consumer performance.

Webpage loading time is between 1 and 3 seconds for most countries using xDSL technology with Ireland as the sole exception (as was the case during the previous testing period), exhibiting a loading time of 3.58 seconds. Loading times are consistently below 1 second for countries making use of Cable and FTTx technologies. Some exceptions include Lithuania and Slovenia for FTTx technology as well as Spain and Malta for cable technology. Poland's loading times also exceed 1 second for both FTTx and cable technologies.

Countries that exhibited excellent actual download throughput also display low webpage loading times, a fact made clear with the lower average loading times for cable and FTTx technologies. However, web browsing speed does not improve proportionately with download throughput. This is because webpage loading time is not just a function of line speed, but also latency, and for services offering 10Mbps download speed or more, latency dominates web browsing performance. This can be seen through the average webpage loading times of cable and FTTx technologies being almost identical. The average loading time for cable based services is slightly higher than FTTx as well, although the difference is negligible. All access technologies have shown a small increase in their average performances since October 2013.

Some exceptions to the above exist where countries with relatively high latency exhibit lower loading times compared to other countries. For example, Portugal, which demonstrated a very high latency for FTTx services also reports one of the lowest loading times. The opposite may also occur, with countries showing lower latency also displaying higher loading times, such as FTTx services in Norway and Bulgaria. This implies the connectivity between our test websites and the Portuguese ISPs is better than to our measurement servers, and the opposite is true in the case of Norway and Bulgaria.

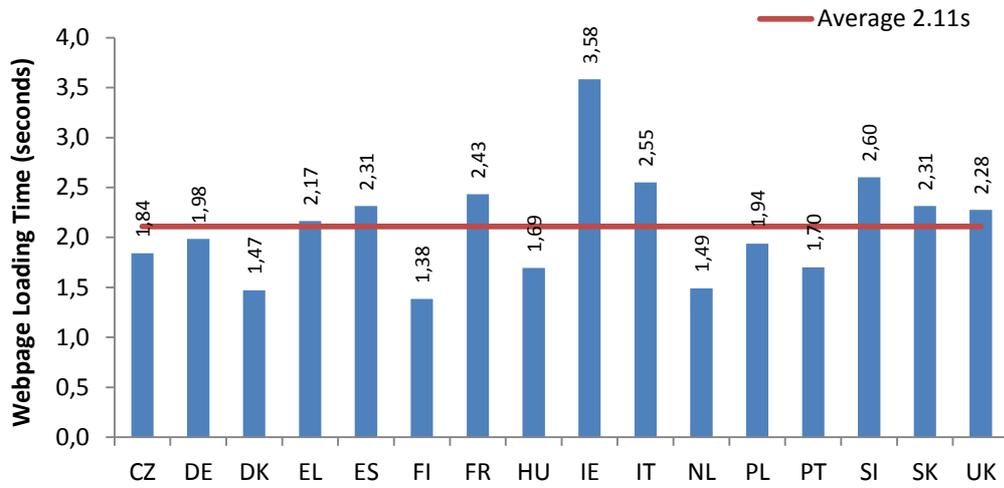


Figure EU.2-25: Webpage Loading Time of xDSL technology during peak periods, by country

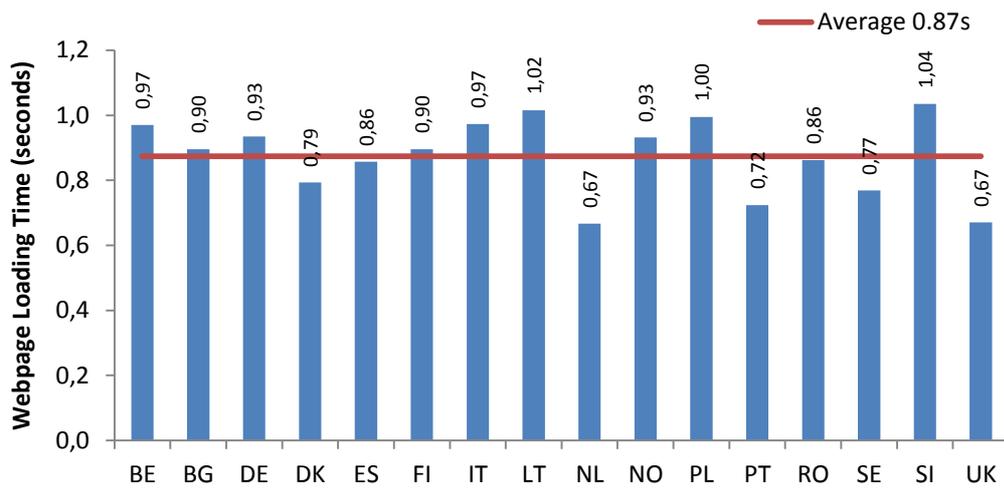


Figure EU.2-26: Webpage Loading Time of FTTx technology during peak periods, by country

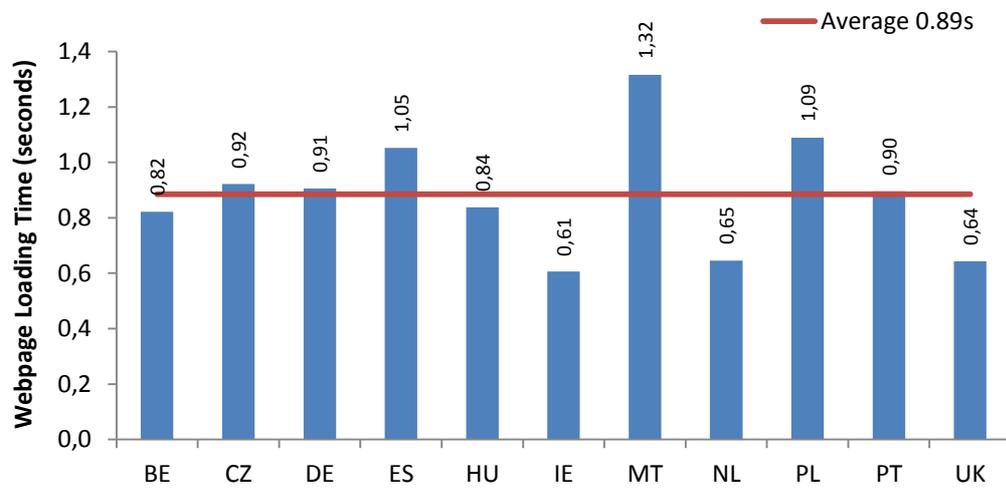


Figure EU.2-27: Webpage Loading Time of cable technology during peak periods, by country

D.2.7 VoIP Jitter

Jitter is an important metric for users who frequently use realtime communication applications. It can also be referred to as latency consistency. Broadband connections frequently shifting between 10ms and 20ms latency would have a high jitter value. This pattern would be very noticeable to consumers using realtime applications such as video streaming and online gaming. Thus, it is better to have lower jitter.

This study reports on downstream and upstream jitter separately. Both are important for two-way communications such as phone calls, but significant technological differences make it so results in the downstream and upstream directions are noticeably divergent.

Downstream jitter is shown in figures EU.2-28 to EU.2-30 during the peak period for each access technology, split by country. Jitter in the downstream direction is very low for most countries across all access technologies. Only xDSL services in a handful of countries show an average significantly above 1ms, such as Ireland, Germany, France, Portugal and Slovenia, with xDSL technology also displaying the highest average compared to cable and FTTx based services. The downstream jitter performance of xDSL services improved greatly in Poland and Slovakia since the previous testing period. Average downstream jitter is slightly higher for FTTx and xDSL based services, in contrast to cable technology which shows a lower average since October 2013. The lower level of jitter for cable technology compared to FTTx technology also reflects the behaviour of latency on each service.

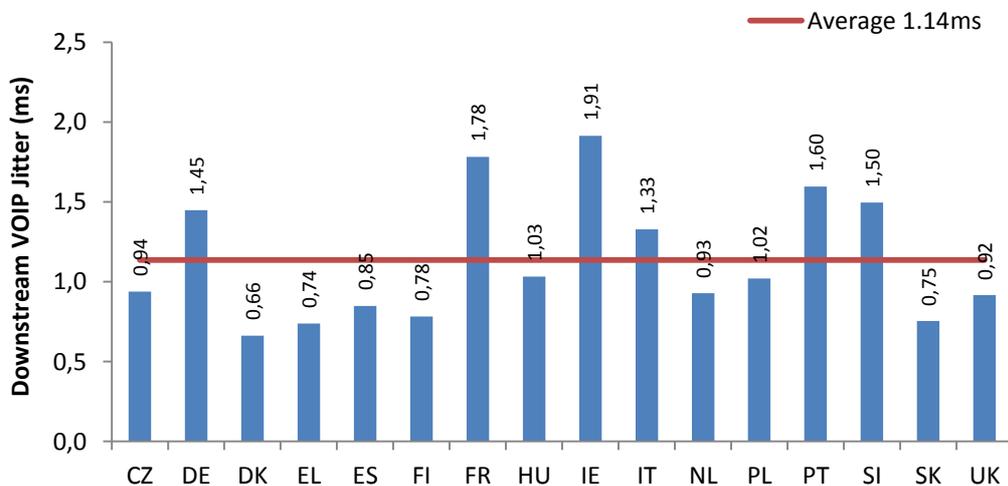


Figure EU.2-28: Downstream VoIP Jitter of xDSL technology during peak periods, by country

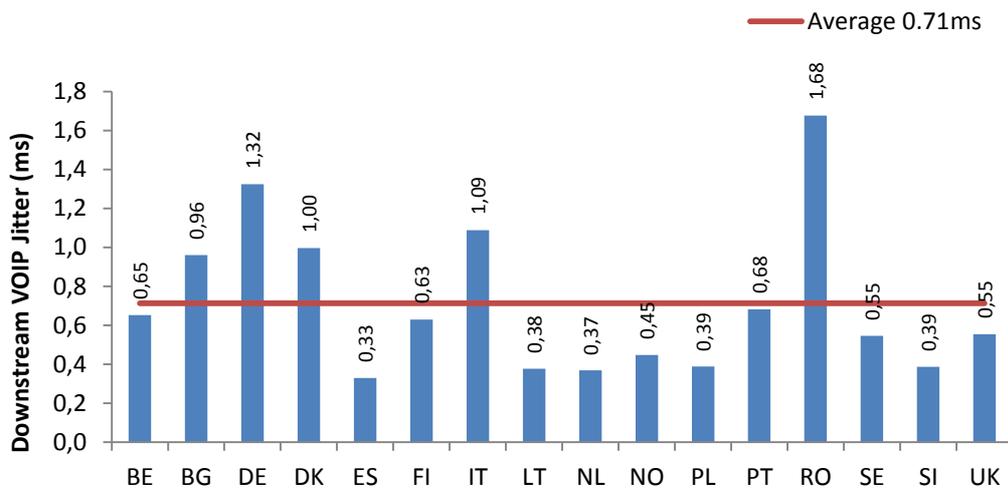


Figure EU.2-29: Downstream VoIP Jitter of FTTx technology during peak periods, by country

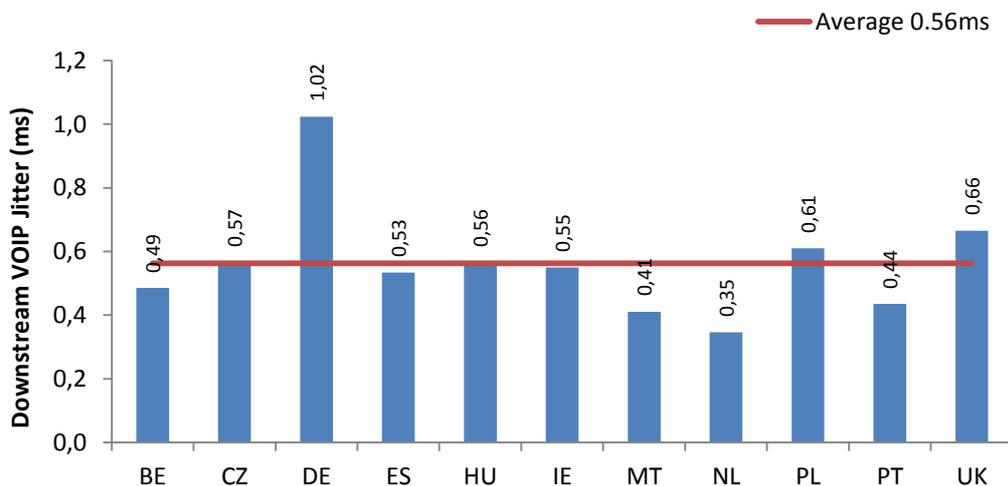


Figure EU.2-30: Downstream VoIP Jitter of cable technology during peak periods, by country

Figures EU.2-31 to EU.2-33 show results for upstream jitter during the peak period for all access technologies, split by country. The xDSL based services in Ireland show upstream jitter is closer to the overall average compared to downstream jitter, with most other nations using xDSL services performing below the average. An exception to this is Hungary, which displays the highest level of upstream jitter among countries offering xDSL services. Germany's upstream jitter for xDSL services improved greatly, but most other nations display higher upstream jitter than in October 2013. This contrasts with FTTx services as upstream jitter rarely exceeded 2ms, with few countries showing upstream jitter higher than 1ms.

Average upstream jitter of cable technology also improved since October 2013, although it still displays the highest average upstream jitter of 3.02ms. Initially, this may appear to be a contradiction with downstream jitter results as well as most other metric performances. The reason for cable's high upstream jitter is due to the

fact that they are based on the concept of TDMA (Time Division Multiple Access). This means the modem's time is divided in slots during which it can either send or receive data, but not both simultaneously. If the modem is busy while the user tries to send a packet, the packet will have to wait in a queue until an opportunity to be transmitted arises. This can result in small but common variations in packet delays.

It is also important to remember that although upstream jitter is relatively high for cable networks, its actual level is low enough that it can be deemed negligible with regards to overall broadband performance. For instance, most Voice over IP (VoIP) phones have a de-jitter buffer of at least 25ms. This means jitter under 25ms would not affect the call quality. As can be seen in figures below, the upstream jitter cable technology across each country is far below 2.5ms.

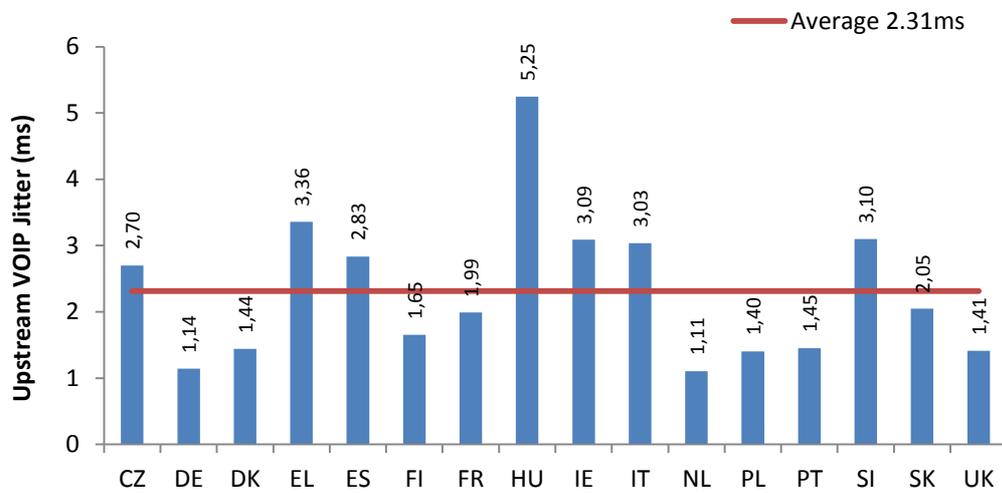


Figure EU.2-31: Upstream VoIP Jitter of xDSL technology during peak periods, split by country

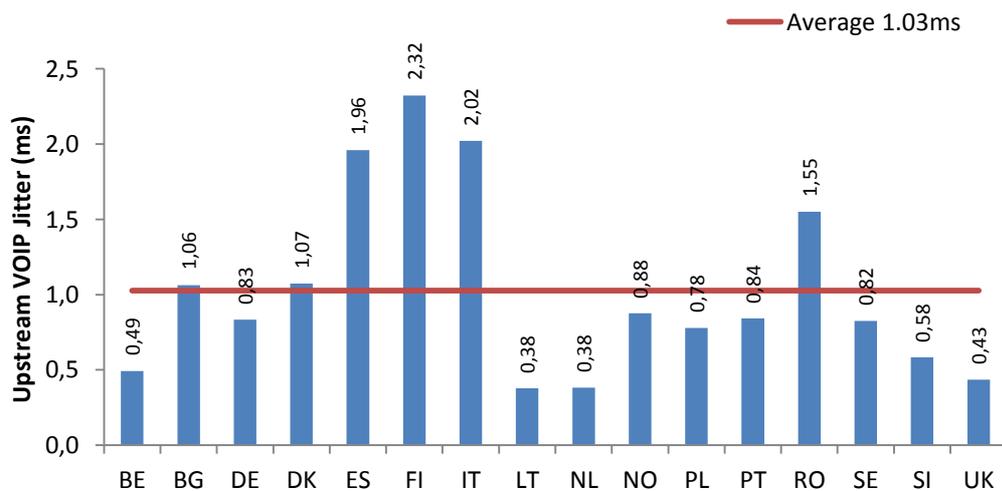


Figure EU.2-32: Upstream VoIP Jitter of FTTx technology during peak periods, by country

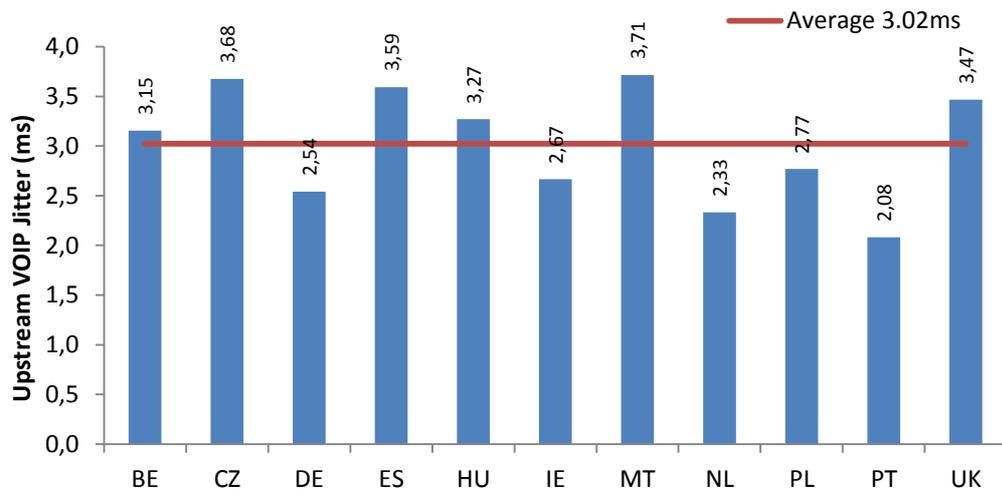


Figure EU.2-33: Upstream VoIP Jitter of cable technology during peak periods, by country

Section E & F Overview

Sections E and F present the same results as sections C and D, but this time using weighted averages. In section E results are weighted by the number of subscribers per technology in each country. In section F results are weighted by the number of subscribers per ISP in each country. This allows us to counteract biases in our sample which may cause us to under-represent or over-represent a given technology or ISP in a country. The weights were calculated using data provided by each NRA in October 2014, although the reference date of the data may be different in some cases. In section F, it is only possible to split by ISP as numerous countries do not provide separate market shares for each ISP by technology, with most ISPs having a market share only within one type of access technology. The sections are included in this study for the purpose of comparing weighted results with the unweighted results within the measurement period of October 2014 found in sections C and D.

All charts in sections E and F present weighted averages with the exception of the CDF (distribution) plots. The CDF plots cannot be weighted, as they present individual measurement results rather than averages (thus there is no average to weight).

The weighted European averages for download and upload are noticeably lower, with latency considerably higher, than the unweighted averages. This is caused by our sample underrepresenting xDSL users, and these users tend to have poorer results than their FTTx and Cable counterparts. Once the data is weighted (effectively skewing in favour of the xDSL results), we find that the overall European averages are pulled down.

In general, each metric split by technology shows very little to no difference between the weighted and unweighted variants. Only two exceptions can be seen in download speed of xDSL technology, with the weighted average in the 24-hour and peak periods is approximately 10% lower than the unweighted average. DNS failure rate of FTTx technology is also shown to decrease during peak hours when weighted, contrary to the unweighted DNS failure rate for this access technology. The reason for this is due to an undersampling of results within the 24-hour period for this metric and technology. It should be noted, however, that failure rates do not change significantly between periods and also do not differ considerably when weighted.

Average results for each technology and metric in Section F also do not show any significant differences from average results in Section D. Very few differences exist for specific countries. Examples include the actual download speed of FTTx technology in Poland is lower than Portugal's and the average throughput of FTTx when weighted, where it is higher when unweighted; actual cable upload speed is more similar between Hungary and Ireland when weighted; packet loss of xDSL technology in Germany, Hungary and Slovenia is much lower as well as packet loss of cable technology in Poland; DNS failure rate of xDSL technology in Spain is higher; FTTx upstream jitter in Norway and Finland is considerably lower when

weighted although Finland's upstream jitter remains significantly above the average; downstream jitter of cable technology in Poland is lower as well.

Certain countries' and ISPs' weighted results will not differ from their unweighted results. This is because there was not sufficient data within the samples to allocate accurate weightings. The countries affected were Greece, Latvia, Lithuania and Portugal. These countries' access technologies and their ISPs were therefore given neutral weightings, leading to weighted results that are no different from their unweighted results.

It should be noted that the weighting methodology described above leads to an over-representation of higher speed tiers in certain countries. This issue will be addressed in more detail in Section G.

E EU Level Analysis - Weighted⁵

E.1 Key Performance Indicators

E.1.1 Download Speed

Figure EU.3-1 and Fig EU.3-2 display actual download speed as a percentage of advertised speed over the 24-hour and peak periods, split by access technology. Cable technology outperforms all other access technology, achieving the highest level of download throughput as a percentage of advertised speed with 85.37% during the peak period, representing a slight drop from results over the 24-hour period. FTTx and xDSL technologies deliver 80.53% and 53.85% of advertised speed respectively during this period. In the case of xDSL technology, this is approximately 10 percentage points than the unweighted result. This is attributable to an undersampling of the access technology. All technologies experienced a small decline in percentage terms during both the peak and 24-hour periods since October 2013. However, this is due to an increase in average headline speeds of each access technology as opposed to a real decline in overall performance. All access technologies apart from xDSL experienced significant improvements in actual throughput, especially cable technology.

Please note that the figures below are not obtained by dividing the average actual speed through the average advertised speed. They are computed on a per-panellist basis and averaged to form an overall figure. This approach is therefore a mean of ratios as opposed to a ratio of means.

	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Technology and Period								
Weighted Actual Speed (Mbps)								
Oct-14	7.83	8.09	75.76	80.34	51.31	52.67	22.92	23.92
Advertised Speed (Mbps)								
Oct-14	15.64	15.64	92.04	92.04	64.52	64.52	32.61	32.61
Weighted Actual/Advertised Speed								
Oct-14	53.85%	55.63%	85.37%	89.70%	80.53%	82.46%	61.72%	63.86%

Figure EU.3-1: Weighted Actual Peak and 24-hour Period Download Speed as a Percentage of Advertised Speed, by technology (higher is better)

⁵ EU refers to the average of all the countries included in the sample, i.e. EU28 countries & Iceland and Norway

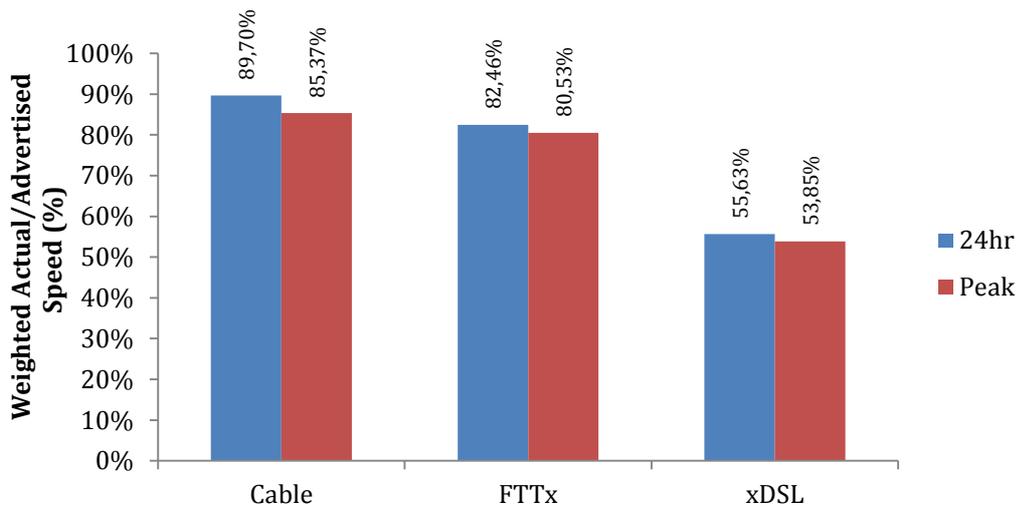


Figure EU.3-2: Weighted Actual Peak and 24-hour Period Download Speed as a Percentage of Advertised Speed, by technology (higher is better)

Figure EU.3-3 below shows actual download speed as a percentage of advertised speed, split by time of day and technology. All technologies exhibit a similar pattern throughout the day, with download speed declining slightly during the day, followed by a more significant drop in throughput during the peak period. xDSL exhibits the smallest amount of decline during peak hours compared to cable and FTTx technologies, with cable technology exhibiting the sharpest drop in throughput during this time. As was shown in figure EU.3-2, cable outperforms all other technologies.

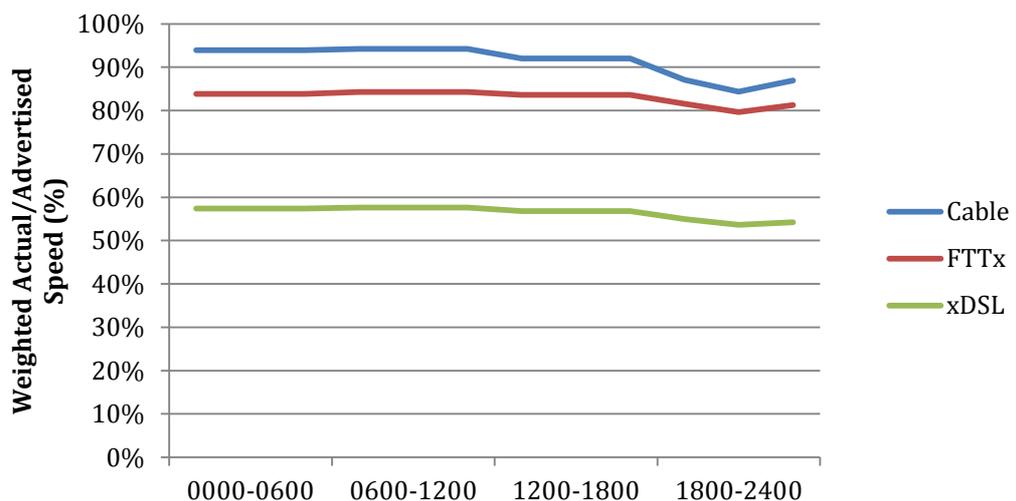


Figure EU.3-3: Weighted Actual Download Speed as a Percentage of Advertised Speed, by hour of day and technology (higher is better)

Figure EU.3-4 below shows actual download speed split by time of day and access technology. As seen in figure EU.3-3, all technologies display similar diurnal patterns. Download speed experiences a small decline in actual throughput in the afternoon followed by a sharper drop in the peak period, as is indicated in figure EU.3-3. xDSL technology shows the smallest amount of change during this period, contrasting with cable technology which experiences the largest decline. Additionally, all technologies apart from xDSL experienced significant improvements in actual download performance since October 2013, particularly cable technology.

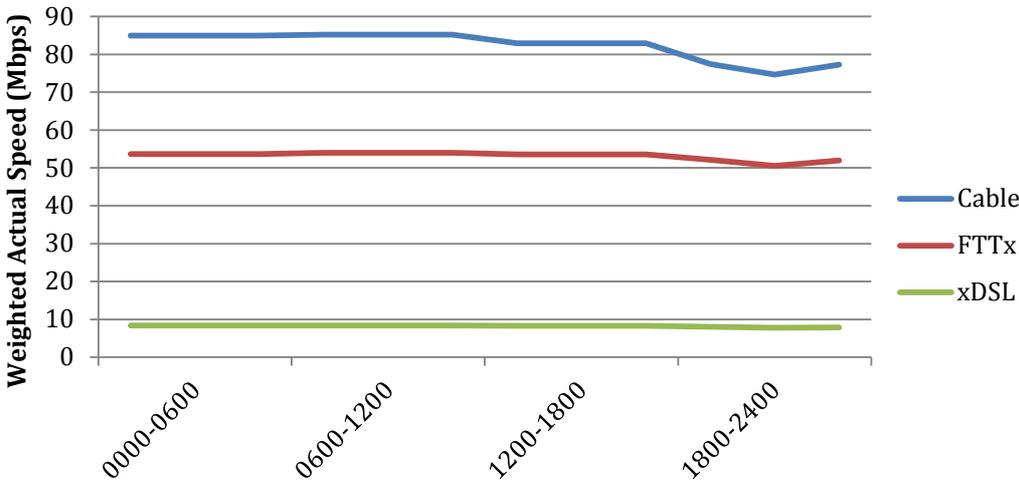


Figure EU.3-4: Weighted Actual Download Speed by hour of day and technology (higher is better)

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

Figure EU.3-5 presents a cumulative distribution chart of the download speed as a percentage of advertised speed, split by access technology. This chart is meant to represent the percentage of consumers who receive at least a certain level of their advertised broadband speed i.e. the chart gives an indication of what proportion of consumers receive their respective advertised speeds. Figures EU.3-2 to EU.3-4 focused on averages alone, but these can mask high levels of inconsistency. The cumulative distribution plot helps to show if there is a significant spread of results within the measurement samples.

For example, one technology may deliver 90% of advertised speed to all its users at all times, and another technology may deliver anything between 60% and 100% of advertised speed. Both may produce an average speed of 90%, thus proving hard to distinguish in charts EU.3-2 to EU.3-4 above. The cumulative distribution chart presents these differences clearly.

Figure EU.3-5 shows 80% of cable consumers receive at least 75% of advertised speed, FTTx consumers receive approximately 70% and xDSL consumers much less than other technologies with 36% of the advertised rate. In the case of cable technology, this appears to be a slight decrease from October 2013. This is due to higher overall advertised rates, with actual throughput seeing an improvement in performance. xDSL and FTTx technologies instead see a slightly greater percentage of consumers receiving better speeds. Results for xDSL consumers show a much wider distribution compared to cable and FTTx, with only 31% of xDSL consumers receiving 80% of advertised speed or better. It is also 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.x.

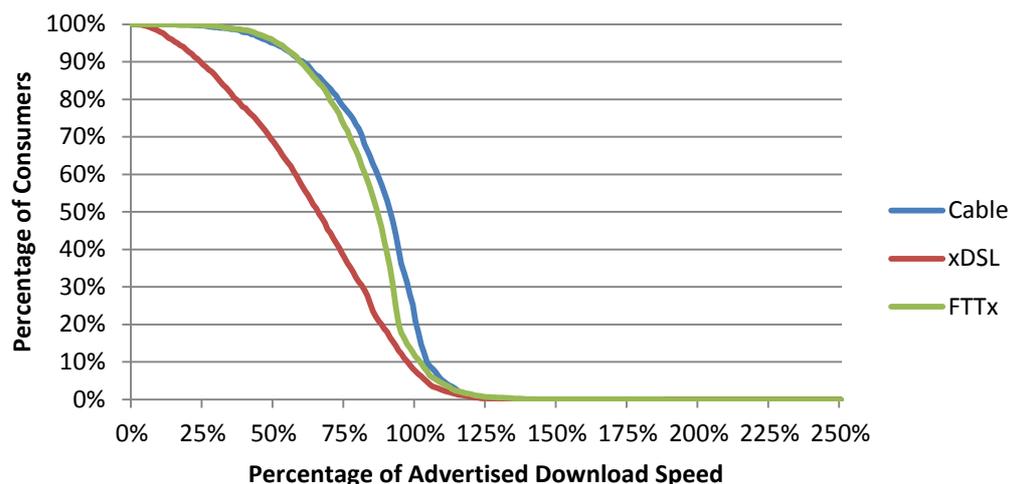


Figure EU.3-5: Cumulative Distribution of Download Speed as a Percentage of Advertised Speed, by technology

E.1.3 Upload Speed

Figures EU.3-6 and EU.3-7 display actual upload speed expressed as a percentage of advertised speed over the peak and 24-hour periods split by access technology. All technologies achieved a level of upload speed above 80% of advertised speed during both periods apart from xDSL during peak hours, with all technologies showing a small decline since October 2013. As was the case in the previous testing period, cable technology continues to exceed the advertised speed during the 24-hour period despite the decline in performance in percentage terms, although it performs just below headline speed during the peak period, with almost no difference in throughput between both periods. Cable again outperforms all other technologies in percentage terms. All technologies show a very small decrease in upload throughput during the peak period.

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 computed on a per-panellist basis and averaged to form one overall figure, as opposed to dividing multiple averages together.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Weighted Actual Speed (Mbps)								
Oct-14	0.77	0.77	8.60	8.66	19.56	19.78	4.34	4.38
Weighted Advertised Speed (Mbps)								
Oct-14	0.99	0.99	8.70	8.70	22.26	22.26	4.87	4.87
Weighted Actual/Advertised Speed								
Oct-14	79.60%	80.16%	99.50%	100.11%	87.97%	88.55%	83.44%	84.01%

Figure EU.3-6: Weighted Actual Peak and 24-hour Period Upload Speed as a Percentage of Advertised Speed, by technology (higher is better)

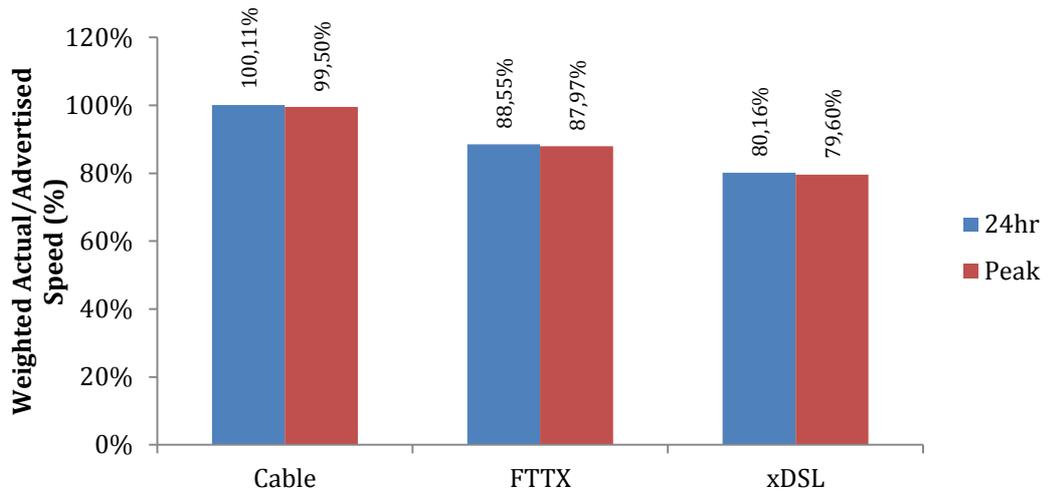


Figure EU.3-7: Weighted Actual Peak and 24-hour Period Upload Speed as a Percentage of Advertised Speed, by technology (higher is better)

Figure EU.3-8 shows actual upload speed as a percentage of advertised speed split by technology and time of day. As was the case in the previous measurement period, Cable and xDSL services deliver a very stable level of throughput during the day, experiencing a very small decline during peak hours. FTTx also displays a relatively stable level of throughput. All technologies show a slightly lower level of throughput in the later hours of the day.

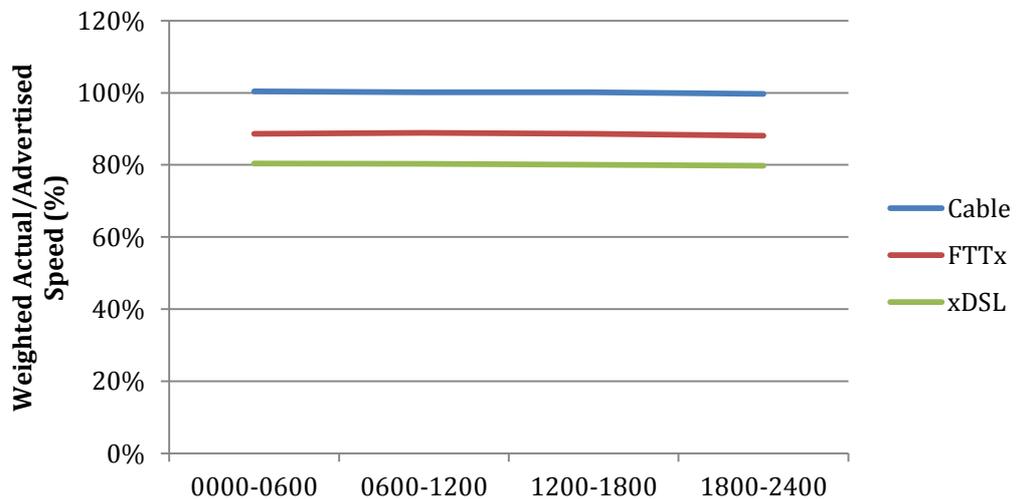


Figure EU.3-8: Weighted Actual Upload Speed as a Percentage of Advertised Speed, by hour of day and technology (higher is better)

Figure EU.3-9 shows actual upload speed by time of day and technology. The behaviour of throughput for all technologies is virtually identical to what is seen in EU.3-8 in the figure below. Cable and xDSL services deliver very stable throughput throughout the day. As was the case in October 2013, a decrease in throughput is more noticeable for FTTx technology throughout the afternoon and peak periods. Contrasting with EU.3-8, FTTx greatly outperforms all other technologies in real terms, indicating much higher advertised rates for FTTx based packages across the countries considered in this report.

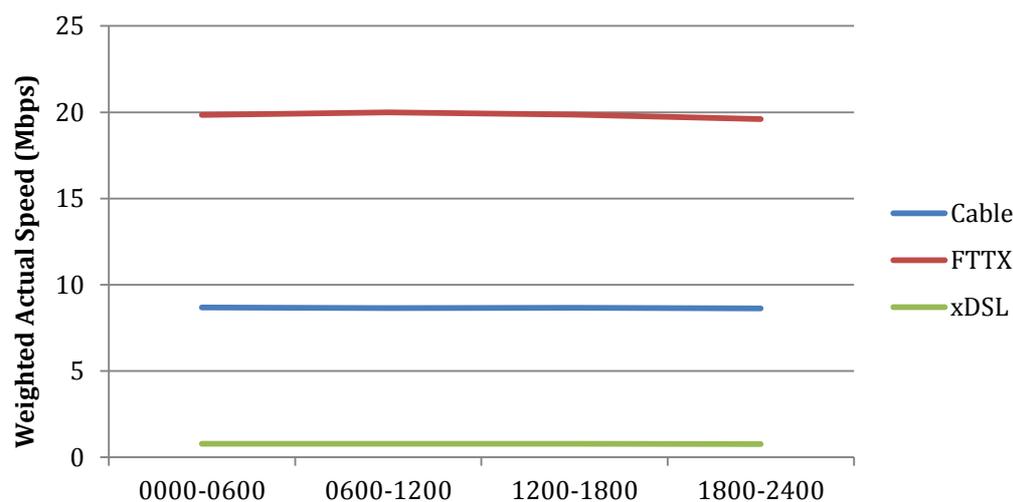


Figure EU.3-9: Weighted Actual Upload Speed by hour of day and technology (higher is better)

Figure EU.3-10 is a cumulative distribution chart for upload speed expressed as a percentage of advertised speed.

The distribution of results is very similar to the previous measurement period, with consumers generally receiving higher levels of advertised speed.

80% of cable consumers are receiving 95% of advertised speed. As was the case in October 2013, FTTx and xDSL are both spread over wider distributions compared to cable technology, with 80% of FTTx and xDSL consumers receiving 80% and 68% of advertised speed respectively. For FTTx technology, this is a significant improvement, in contrast to xDSL technology whose performance is virtually unchanged.

The distribution of results for xDSL consumers is much tighter for upload speed compared to download speed. This is caused by the asymmetric nature of broadband services, with upload speed provisioned at far lower rates.

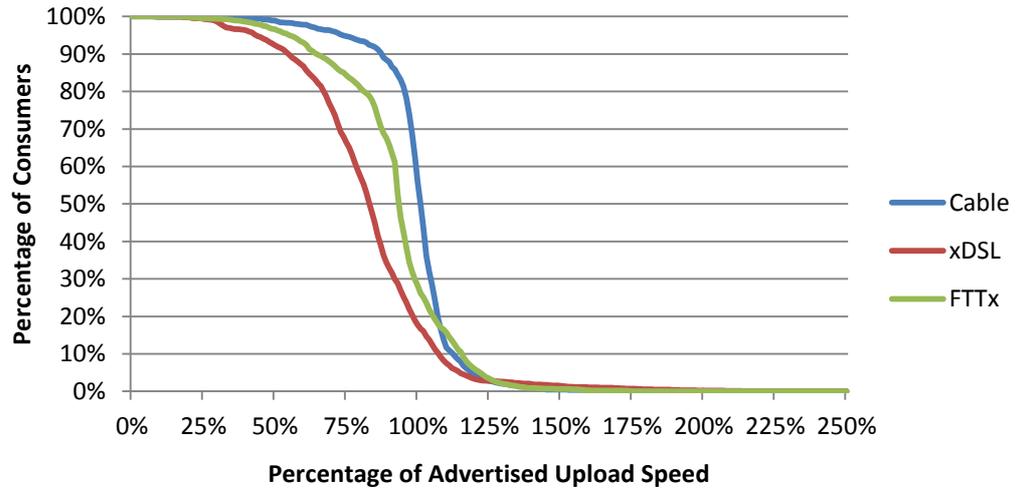


Figure EU.3-10: Cumulative Distribution of Upload Speed as a Percentage of Advertised Speed, by technology

E.1.4 Latency

Figures EU.3-11 and EU.3-12 show average latency over the peak and 24-hour periods across all types of access technology included in this report. xDSL technology displays a slight increase in latency since October 2013, in contrast to cable and FTTx technology which exhibit a slight improvement. Average latency across the EU is noticeably higher, although this is due to the weighting which skews the result further toward xDSL technology, which is much more common than other access technologies. Latency is lower during the 24-hour period compared to the peak period, with xDSL displaying the largest change in real terms.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Weighted Latency (ms)	37.84	35.41	16.35	15.31	18.81	17.66	32.39	30.31

Figure EU.3-11: Weighted Peak period and 24-hour Average Latency results, by technology (lower is better)

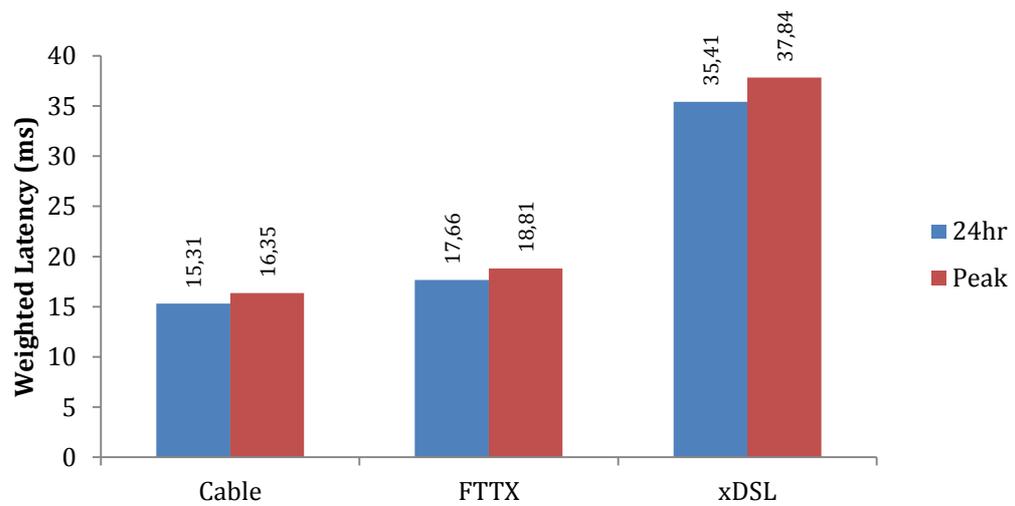


Figure EU.3-12: Weighted Peak period and 24-hour Latency, by technology (lower is better)

Figure EU.3-13 displays latency split by hour of day and technology. As mentioned above for figure EU.3-12, xDSL experiences the most noticeable increase in latency during the peak period. This increase is also more noticeable compared to the previous testing period. In contrast, FTTx latency has become more stable since October 2013.

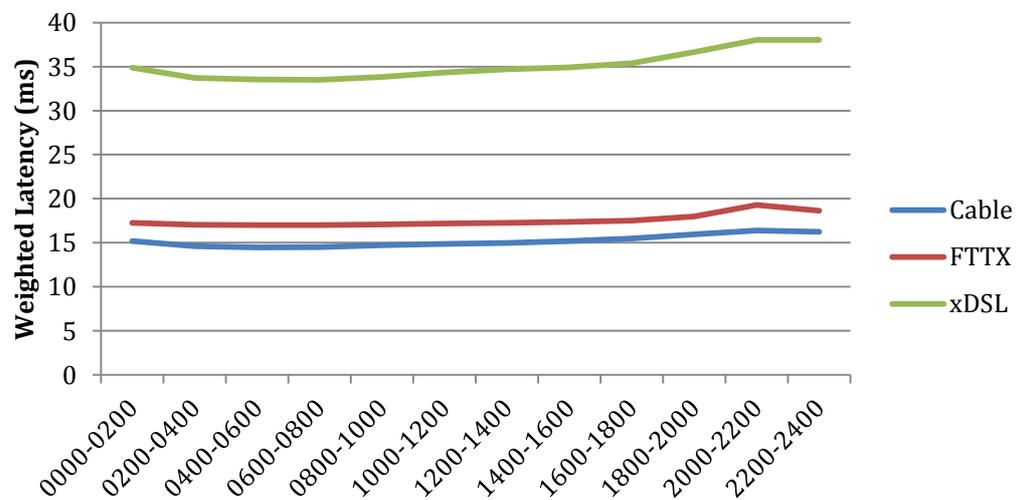


Figure EU.3-13: Weighted Latency by hour of day and technology (lower is better)

Figure EU.3-14 depicts the cumulative distribution of latency. This differs from the CDF charts shown previously for download and upload throughput as it does not display the percentage of advertised figures. This is because ISPs generally do not advertise latency for their broadband products (aside from exceptional cases), as latency can vary wildly depending on the host the user is communicating with. This CDF chart and all others that follow it will instead show the actual value unless stated otherwise.

Figure EU.3-14 shows cable and FTTx technologies have similar distributions of latency across its consumers, as was the case in October 2013. xDSL technology again exhibits a comparatively wider distribution, reflecting the behaviour of download throughput. The CDF plot below also shows that latency during this report's testing period is very close to what it was in October 2013, with 60% of cable and FTTx consumers experiencing latency of up to 17ms and 20ms respectively. It is expected behavior for cable and FTTx to experience lower latency compared to xDSL technology as they are less affected by the length of the "last-mile" cable than xDSL, with 60% of xDSL consumers experience 37ms or less. However, this is a significant improvement for xDSL technology, with 60% of consumers previously receiving 50ms or less.

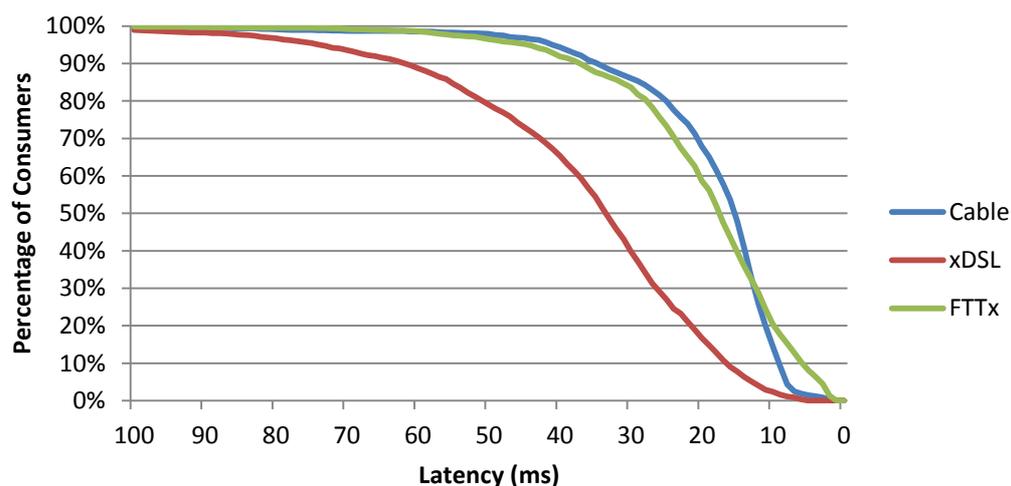


Figure EU.3-14: Cumulative Distribution of Latency, by technology

E.1.5 Packet Loss

Figure EU.3-15 and Figure EU.3-16 show the average packet loss over the peak and 24-hour periods split by access technology. xDSL exhibits the largest level of packet loss across both time periods compared to all other technologies as well as a significantly larger change between the peak and 24-hour periods. In contrast, cable and FTTx experience a very small increase in packet loss during peak hours. Packet loss for xDSL and FTTx technology is lower during this testing period compared to October 2013, with FTTx exhibiting a significant improvement.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Weighted Packet Loss (%)								
Oct-14	0.46%	0.28%	0.20%	0.18%	0.16%	0.14%	0.39%	0.25%

Figure EU.3-15: Weighted Peak period and 24-hour packet loss, by technology (lower is better)

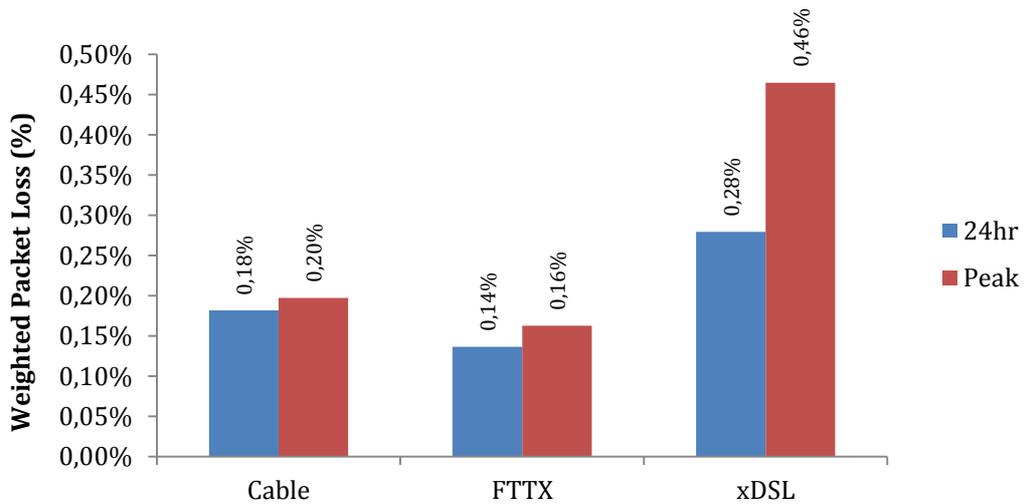


Figure EU.3-16: Weighted Peak period and 24-hour packet loss, by technology (lower is better)

Figure EU.3-17, which displays packet loss by hour of day and split by technology, demonstrates the noticeable increase in packet loss for xDSL services during the day, particularly in the peak period. All types of technology also behave very similarly throughout the day, slightly decreasing in the morning period followed by an increase during the late afternoon and another, sharper increase during the peak hours. These changes are more significant for packet loss of xDSL technology than for cable and FTTx. Packet loss for xDSL technology is also consistently higher than it is for cable and FTTx throughout the day.

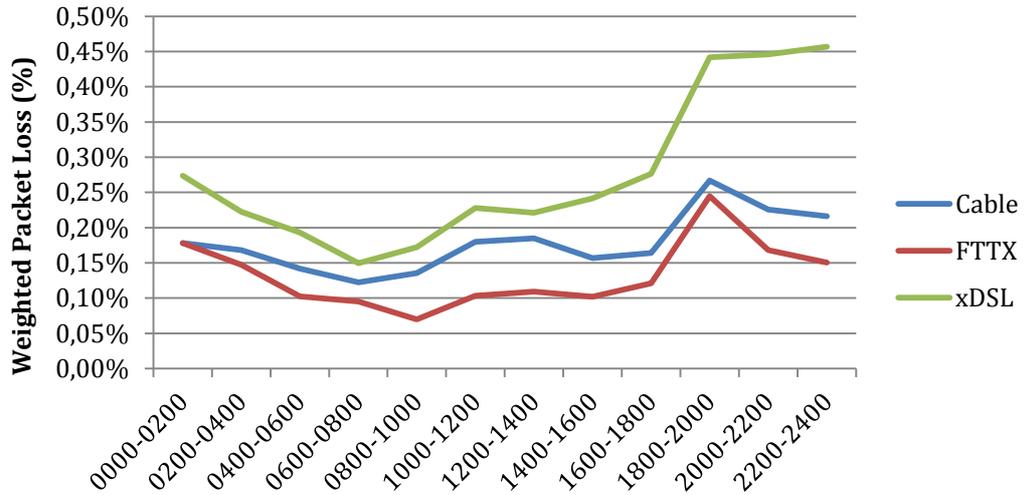


Figure EU.3-17: Weighted Packet Loss by hour of day and technology (lower is better)

Figure EU.3-18 is the cumulative distribution plot for packet loss. As was the case in October 2013, very few consumers experience high levels of packet loss. 90% of xDSL consumers experience only 1% or less packet loss. It should also be noted that xDSL has a slightly wider distribution than cable and FTTx, whose distributions are virtually identical. This is a strong indication that the vast majority of consumers experience very low packet loss.

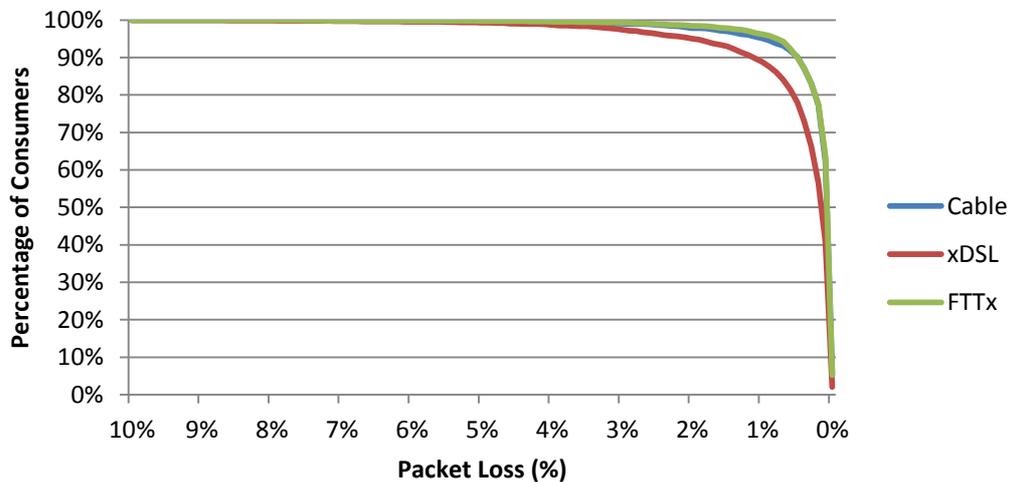


Figure EU.3-18: Cumulative Distribution of Packet Loss, by technology

E.1.6 **DNS Resolution and Failure Rate**

DNS Resolution Time

Figures EU.3-19 and EU.3-20 shows DNS resolution time during the peak and 24-hour periods. As with latency and packet loss, DNS resolution time is much higher for xDSL based services compared to cable and FTTx. This is to be expected as DNS is directly affected by the round-trip latency of the underlying technology. All access technologies apart from FTTx exhibit lower DNS response times than in October 2013. FTTx technology instead shows slightly higher response times, although it also shows the smallest change between the peak and 24-hour periods.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Weighted DNS Resolution (ms)								
Oct-14	36.23	35.34	16.61	15.84	19.09	18.88	31.18	30.4

Figure EU.3-19: Peak Period and 24-hour DNS Resolution Time, split by technology (lower is better)

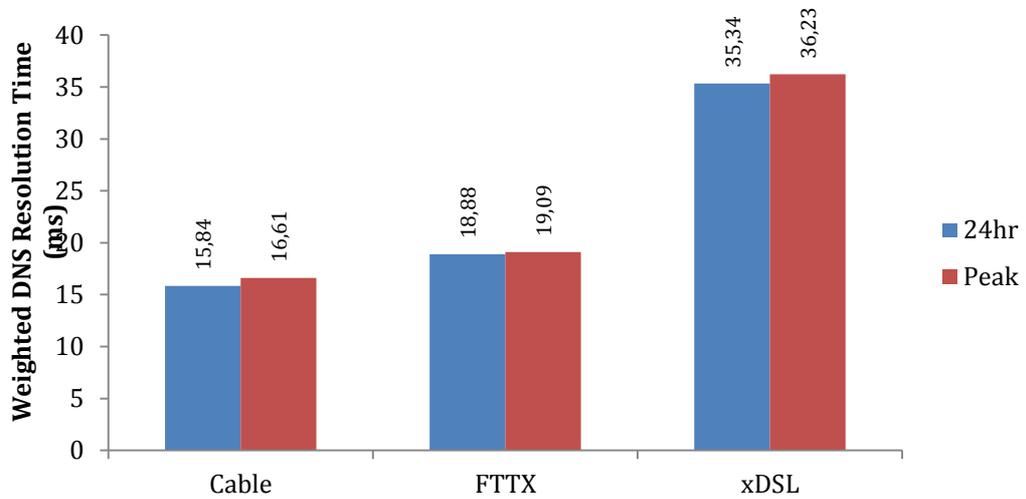


Figure EU.3-20: Weighted Peak Period and 24-hour DNS Resolution Time, split by technology (lower is better)

Figure EU.3-21 below displays DNS resolution time split by time of day and access technology. As indicated in figure EU.3-20, FTTx remains very consistent throughout the day, experiencing almost no change during peak hours. DNS resolution times for cable are lower than for FTTx even following an increase during the afternoon and peak periods. xDSL also increases steadily throughout the day and slightly more during peak hours.

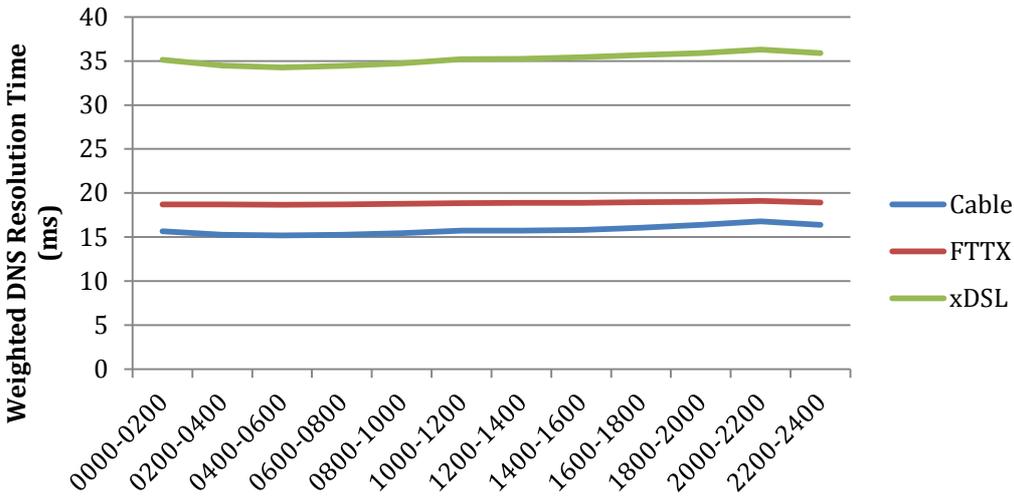


Figure EU.3-21: Weighted DNS Resolution Time, split by hour of day and technology (lower is better)

Figure EU.3-22 shows the cumulative distribution plot for DNS resolution split by access technology. The figure below shows virtually no change in DNS resolution from the previous testing period, with 50% of cable and FTTx consumers seeing DNS resolution times of 15ms or less, as was the case in October 2013. DNS resolution is much higher for xDSL users, as was shown above in figures EU.1-20 and EU.1-21 as well.

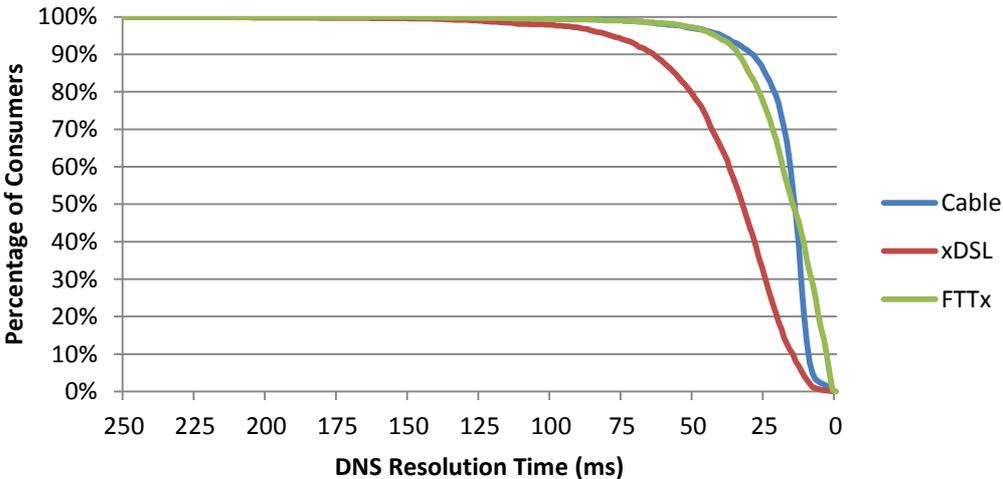


Figure EU.3-22: Cumulative Distribution of DNS Resolution Time, by technology

DNS Resolution Failure Rate

Figures EU.3-23 and EU.3-24 show peak period and 24-hour DNS resolution failure rate across all access technologies. Cable technology displays significant improvements in failure rates since October 2013, showing much lower figures during the peak and 24-hour measurement periods as well as virtually no change between the two periods. Failure rates are slightly higher during the peak period only for xDSL technology. In contrast, FTTx technology shows a lower failure rate during peak hours than the 24-hour period, although this is more likely due to a slight undersampling of peak period results compared to 24-hour period figures for this access technology, disproportionately influencing the weightings for each measurement period. xDSL technology also experiences an improvement since October 2013, particularly during the peak period.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
DNS Failure Rate (%)								
Oct-14	0.36%	0.30%	0.19%	0.19%	0.18%	0.22%	0.31%	0.27%

Figure EU.3-23: Peak period and 24-hour DNS Resolution Failure Rate, split by technology (lower is better)

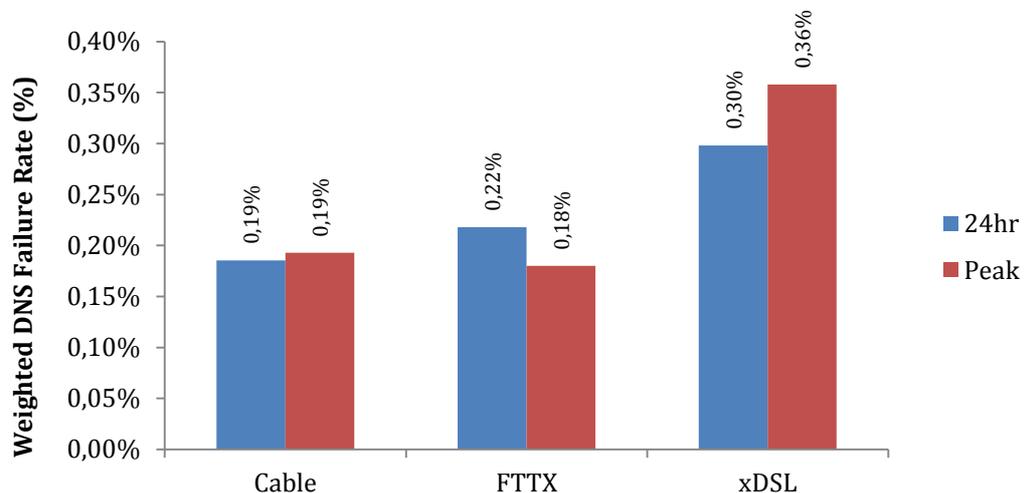


Figure EU.3-24: Weighted Peak period and 24-hour DNS Resolution Failure Rate, by technology (lower is better)

Figure EU.3-25 shows DNS resolution failure rate by time of day and split by access technology. This figure shows that the behaviour of failure rates of cable & FTTx technologies is very uneven for all technologies, particularly during the early morning period. All access technologies display slightly higher failure rates during the afternoon and peak periods, especially xDSL technology. FTTx technology is an exception to this, showing a much sharper increase during the morning and afternoon periods, becoming lower and more stable during peak hours.

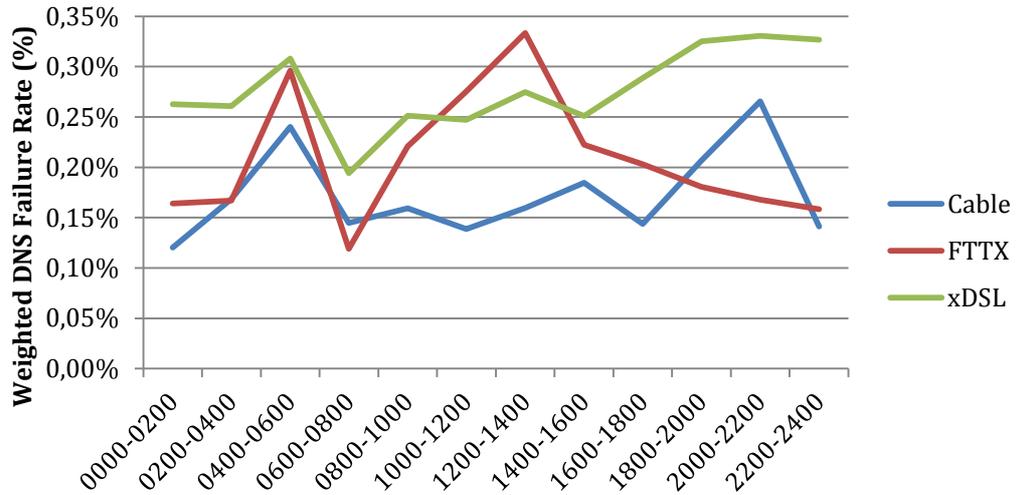


Figure EU.3-25: Weighted DNS Resolution Failure Rate, by hour of day and technology (lower is better)

Figure EU.3-26 shows the cumulative distribution chart of DNS resolution failure rate from each type of access technology. As was the case with packet loss, the chart below shows that high failure rates are uncommon among the majority of users for each type of technology, with approximately 90% of users of all access technologies seeing less than 1% failure rates.

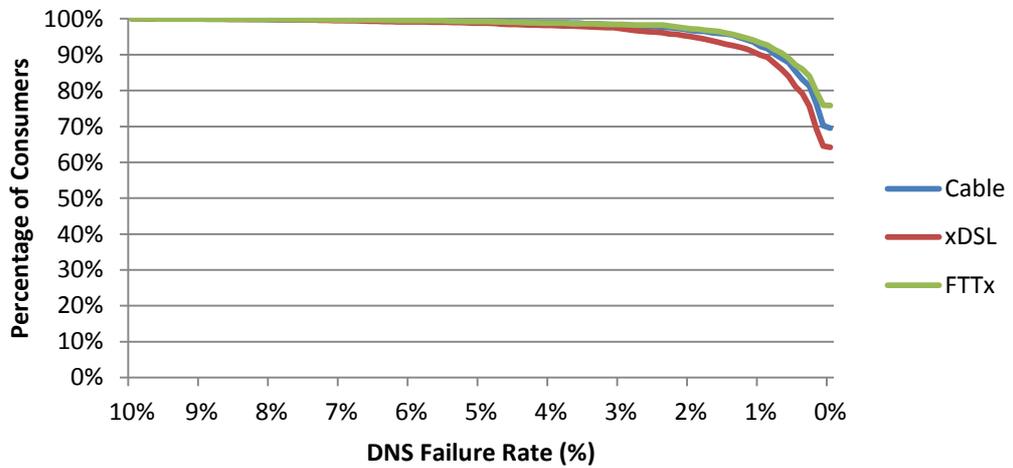


Figure EU.3-26: DNS Resolution Failure Rate, by hour of day and technology (lower is better)

E.1.7

Web Browsing

Figures EU.3-27 and EU.3-28 show the average webpage loading times across all the access technologies covered in this study during the peak and 24-hour periods. It should be noted that this test measures against real websites e.g. 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 exhibits by far the highest webpage loading times during both the peak and 24-hour periods. Web browsing speed for xDSL technology proved to be more than twice as slow as cable and FTTx services, which performed nearly identically during the 24-hour and peak periods, with FTTx proving to be slightly quicker. FTTx services also experienced the smallest increase in loading times during the peak period. Webpage loading times for all access technologies is higher during this report’s measurement period compared to October 2013, which may be due to larger page sizes and additional objects on each page.

The behaviour of webpage loading times also mimics the patterns exhibited by latency and DNS resolution time, which is to be expected given that web browsing performance is a function of both line speed and round trip latency. Additionally, beyond a certain level of downstream throughput, usually 10Mbps, latency becomes the main element affecting web browsing performance. This is shown by results in the figures below.

Technology and Period	xDSL Peak	xDSL 24hr	Cable Peak	Cable 24hr	FTTx Peak	FTTx 24hr	EU Peak	EU 24hr
Webpage Loading Time (s)								
Oct-14	2.28	2.13	0.74	0.69	0.74	0.7	1.86	1.74

Figure EU.3-27: Weighted Peak period and 24-hour Webpage Loading Time, by technology (lower is better)

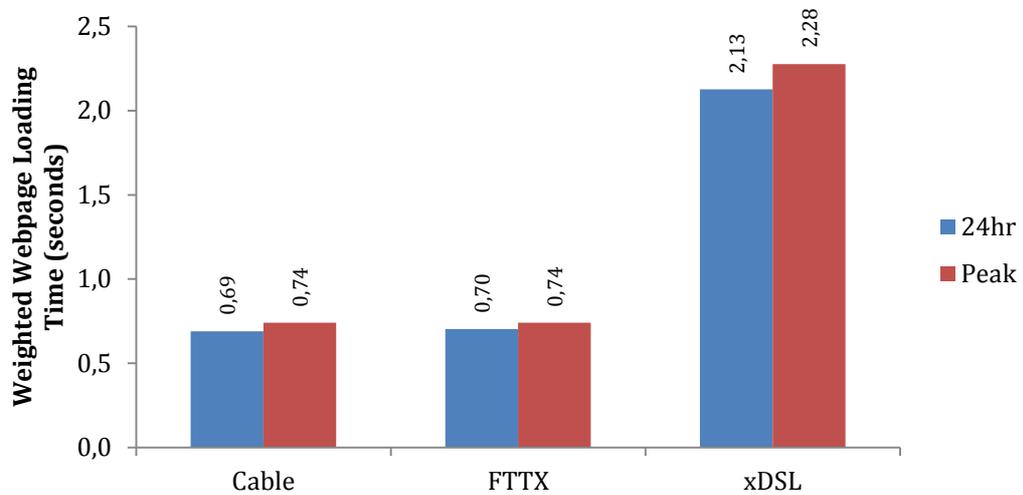


Figure EU.3-28: Weighted Peak period and 24-hour Webpage Loading Time, by technology (lower is better)

Figure EU.3-29 below shows the average webpage loading time by time of day and access technology. All technologies exhibit very similar patterns to latency throughout the day. As was the case in October 2013, FTTx technology displays the most consistent loading times, exhibiting the smallest increase during peak hours compared to other access technologies, although cable technology performs almost identically. Peak hour changes are more evident for xDSL technology, which also shows much higher webpage loading times.

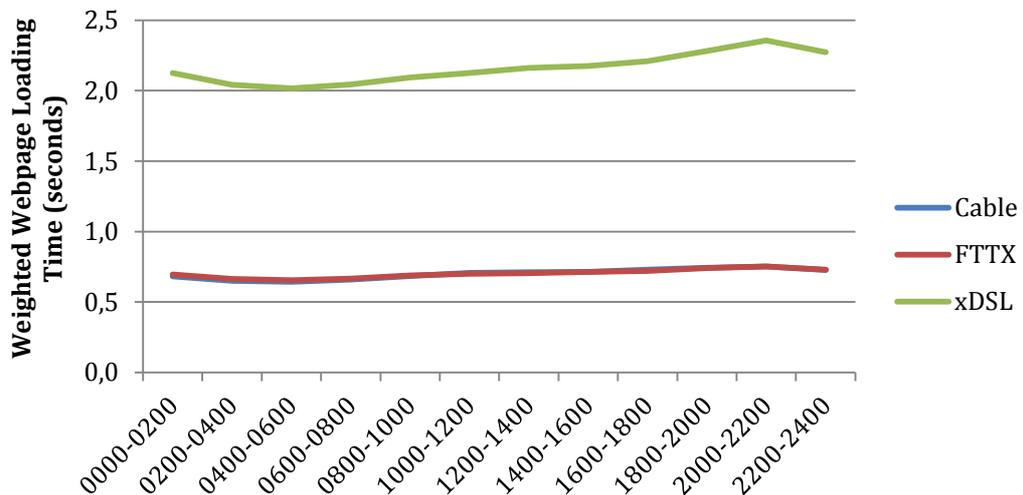


Figure EU.3-29: Weighted Webpage Loading Time, by hour of day and technology (lower is better)

Figure EU.3-30 displays the cumulative distribution plot for webpage loading time split by technology. 90% of cable and FTTx consumers experience loading times of 1.2 seconds at most, slightly higher than in October 2013, with both types of technologies exhibiting virtually identical distribution across their respective user bases. 90% of users of xDSL technology, which has a wider distribution than cable and FTTx technologies, see loading times of approximately 3.8 seconds or less, also higher than in the previous year.

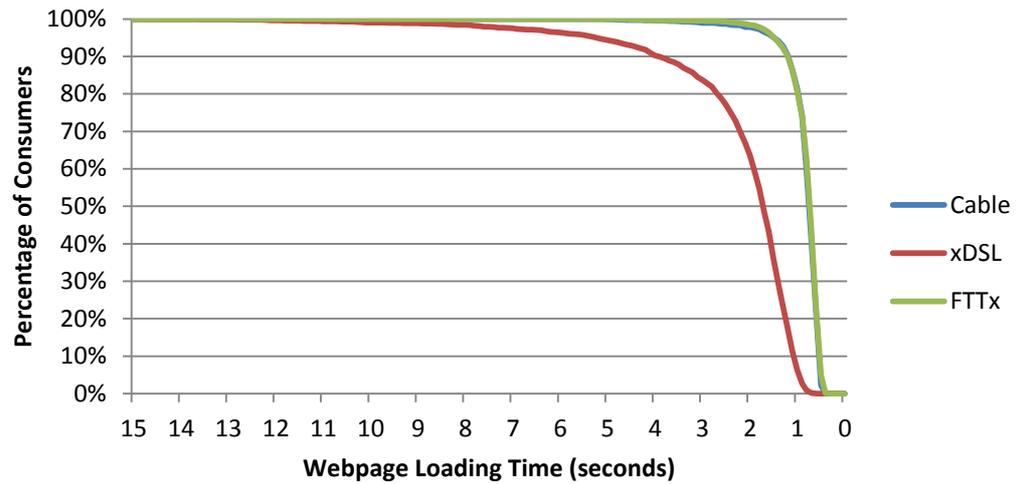


Figure EU.3-30: Cumulative Distribution of Webpage Loading Time, by technology

E.1.8 VOIP

Downstream VoIP Jitter

Figures EU.3-31 and EU.3-32 show downstream jitter during the peak and 24-hour periods across all types of access technology. Downstream jitter has improved slightly for cable technology since October 2013, displaying lower jitter during the 24-hour period. All other technologies exhibit a slightly higher level of jitter, particularly xDSL technology which also experiences the largest increase in downstream jitter during peak hours.

Technology and Period	xDSL	xDSL	Cable	Cable	FTTx	FTTx	EU	EU
	Peak	24hr	Peak	24hr	Peak	24hr	Peak	24hr
Weighted Downstream Jitter (ms)								
Oct-14	1.18	0.79	0.54	0.35	0.7	0.46	1.03	0.68

Figure EU.3-31: Weighted Peak period and 24-hour Downstream VoIP Jitter, by technology (lower is better)

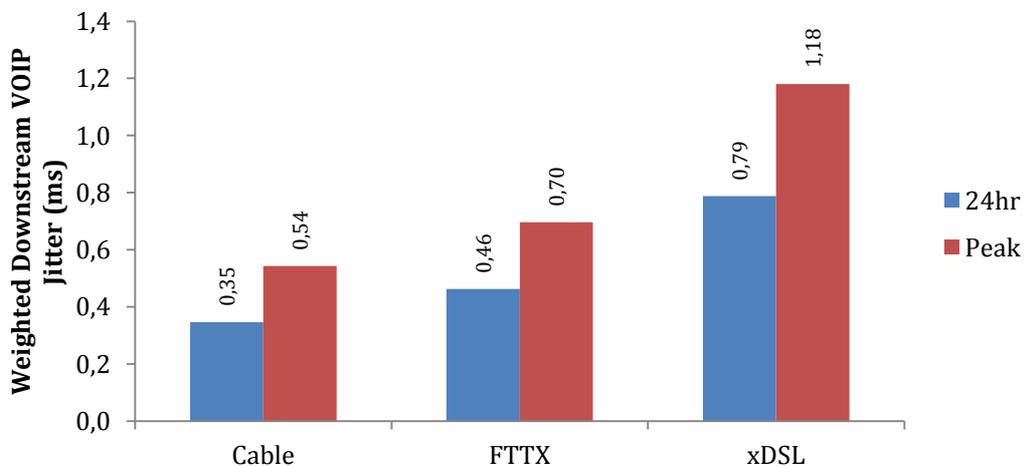


Figure EU.3-32: Weighted Peak period and 24-hour Downstream VoIP Jitter, by technology (lower is better)

Figure EU.3-33 shows downstream jitter split by hour of day and access technology. The behaviour of jitter for all technologies is very similar, increasing steadily throughout the day followed by sharper increases during the peak period. Cable technology displays the lowest level of downstream jitter compared to other access technologies, contrasting with downstream jitter for xDSL services, which is noticeably higher and increases more sharply.

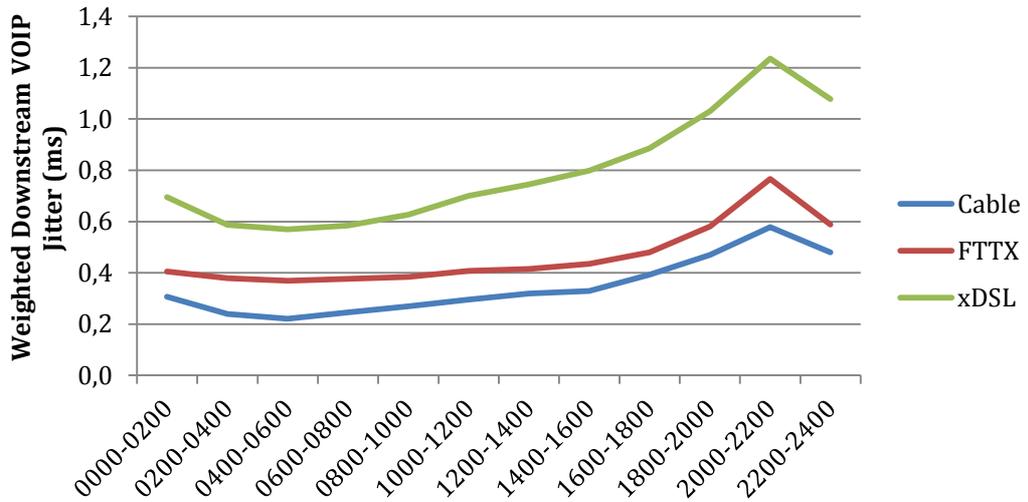


Figure EU.3-33: Weighted Downstream VoIP Jitter, by hour of day and technology (lower is better)

The cumulative distribution for downstream VoIP jitter, shown below in figure EU.3-34, shows a tight distribution across all access technologies covered in this study, particularly cable technology, as was the case for web browsing speed. In tandem with average performance, the distribution for xDSL and FTTx technology also appears to have widened slightly since October 2013. 80% of cable and FTTx consumers experience downstream jitter of 0.7ms and 1ms respectively or less, with 80% of xDSL users seeing jitter of 1.5ms or less.

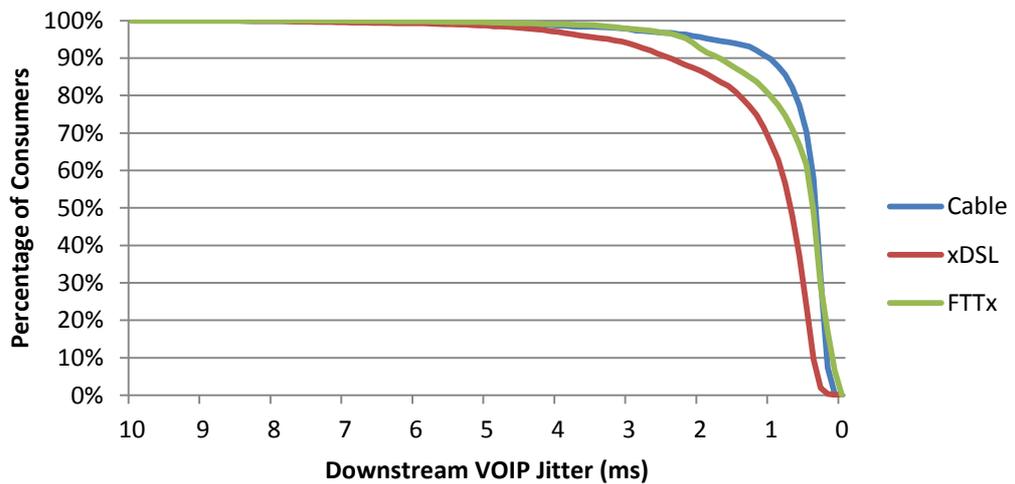


Figure EU.3-34: Cumulative Distribution of Downstream VoIP Jitter, by technology

Upstream VoIP Jitter

Figures EU.3-35 and EU.3-36 show upstream VoIP jitter for the access technologies covered in this study during the peak and 24-hour periods. As was the case in October 2013, cable products exhibit the highest level of upstream jitter across all access technologies and the greatest increase during peak hours. FTTx services never exceeded 1ms of upstream jitter, and also experienced the smallest amount of increase during the peak period. Although the upstream jitter of cable technology is much higher than that of all other access technologies, it is also the only one to display a slight improvement in average performance. FTTx and xDSL technologies instead exhibit a slightly higher level of upstream jitter since the previous year.

The reason for cable services exhibiting higher upstream jitter 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. If the modem is busy whilst the user tries to send a packet, 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 dejitter buffer of at least 25ms, meaning jitter under 25ms would not affect 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	EU Peak	EU 24hr
Weighted Upstream Jitter (ms)								
Oct-14	1.83	1.59	3.01	2.51	0.68	0.61	1.84	1.59

Figure EU.3-35: Weighted Peak period and 24-hour Upstream VoIP Jitter, by technology (lower is better)

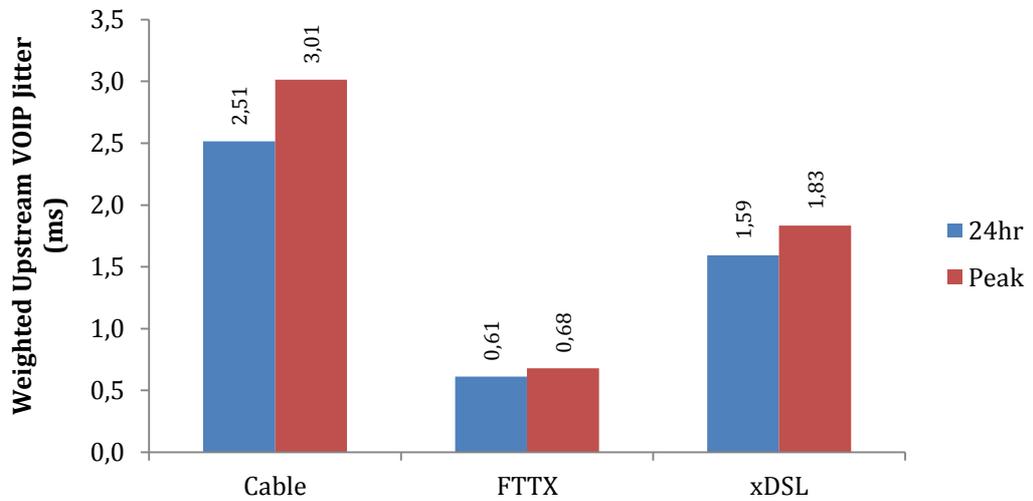


Figure EU.3-36: Weighted Peak period and 24-hour Upstream VoIP Jitter, by technology (lower is better)

Figure EU.3-37, which depicts upstream jitter by hour of day and technology, shows that the behaviour of jitter across all access technologies does not differ significantly from the previous measurement period. Upstream jitter for FTTX and xDSL remains mostly consistent, slightly increasing throughout the day and increasing more sharply during the peak period. This is more evident in xDSL technology. Cable services exhibit a much higher and more severe increase in jitter.

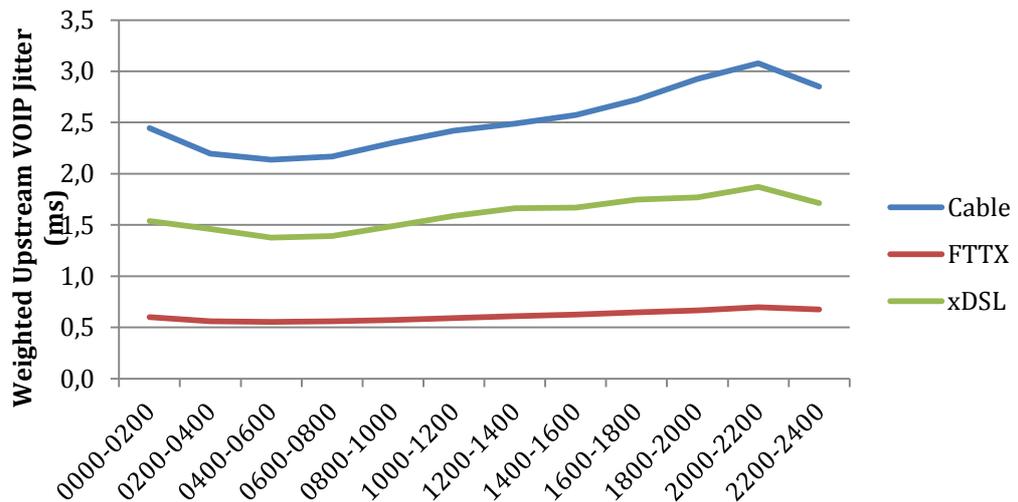


Figure EU.3-37: Weighted Upstream VoIP Jitter, split by hour of day and technology (lower is better)

Figure EU.3-38 shows the cumulative distribution of upstream jitter for all access technologies. The distribution of figures is again very similar to what was observed in October 2013. All technologies, particularly FTTx and xDSL, have a tight distribution. Cable services also display a tighter distribution compared to the last measurement period. 80% of cable users experience 3.6ms or less of upstream jitter, as opposed to 4ms in the previous year. 80% FTTx and xDSL see approximately 1ms and 2ms of jitter respectively, as was the case in October 2013.

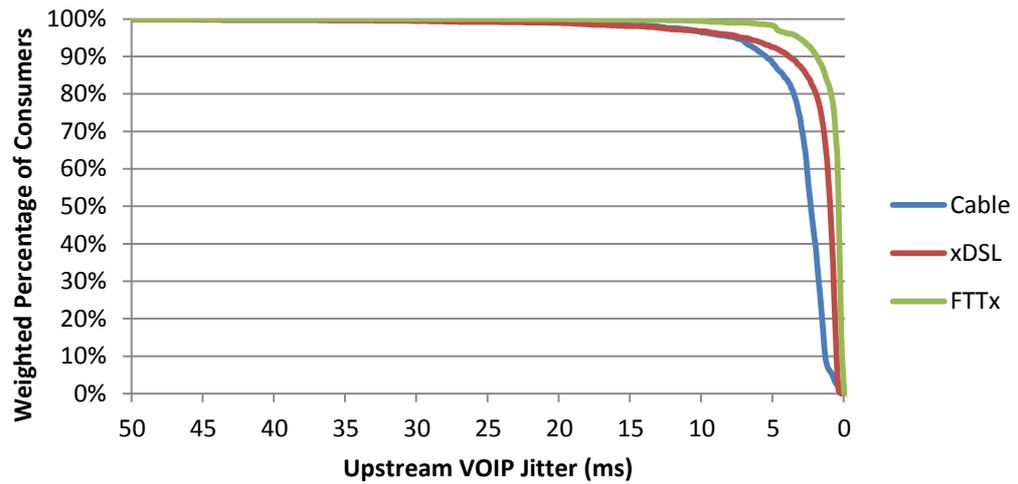


Figure EU.3-38: Weighted Cumulative Distribution of Upstream VoIP Jitter, by technology

E.2 Comparison with the United States⁶

E.2.1 Download Speed

Figure EU.3-39 below shows the average actual and advertised download speeds in Europe and the US for all access technologies. While throughput expressed as a percentage of advertised speed is lower, actual download throughput in Europe is typically better than in the USA across all access technologies other than xDSL technology. This suggests that significant differences in the way products are advertised to the general public exist between the two regions.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual/Advertised (%)	US Advertised (Mbps)	US Actual (Mbps)	US Actual/Advertised (%)
xDSL	15.64	7.83	53.85%	8.41	7.67	91.19%
FTTx	64.52	51.31	80.53%	36.53	41.35	113.19%
Cable	92.04	75.76	85.37%	24.88	25.48	102.42%

Figure EU.3-39: Comparison of Actual and Advertised Download Speed between Europe and the USA, by technology

E.2.2 Upload Speed

Figure EU.3-40 shows the actual and advertised upload throughput for all access technologies in Europe and the USA during the peak period. As was the case in the previous testing period, the USA exhibits better actual throughput performance for xDSL and FTTx technologies compared to Europe. In contrast, cable technology performance much better in Europe. As was the case with download speed, the USA shows much better throughput expressed as a percentage of advertised speed, with cable and FTTx technology exceeding the advertised speed. This again indicates differences in marketing schemes between the two regions.

Technology	Europe Advertised (Mbps)	Europe Actual (Mbps)	Europe Actual/Advertised (%)	US Advertised (Mbps)	US Actual (Mbps)	US Actual/Advertised (%)
xDSL	0.99	0.77	80.16%	0.88	0.86	98.2%
FTTx	22.26	19.56	87.97%	22.79	25.93	113.79%
Cable	8.70	8.70	100.11%	4.07	4.51	110.71%

Figure EU.3-40: Comparison of Actual and Advertised Upload Speed between Europe and the USA, by technology

⁶ Data taken from Measuring Broadband America – April 2014 - <http://www.fcc.gov/measuring-broadband-america>

E.2.3 Latency

Figure EU.3-41 shows a comparison of latency figures between Europe and the USA, split by access technology. Contrasting with results from the previous measurement period, Europe displays much better latency across all access technologies. There is a significant difference in latency for both cable and xDSL, but not for FTTx. This again suggests xDSL and cable based services are deployed in a similar way throughout each region.

Technology	Europe	US
xDSL (ms)	37.84	48.96
FTTx (ms)	18.81	24.07
Cable (ms)	16.35	31.77

Figure EU.3-41: Comparison of Latency between Europe and the USA, by technology

E.2.4 Packet Loss

Figure EU.3-42 is the comparison of packet loss during the peak period between Europe and the USA, split by technology. All access technologies in the USA displayed significantly lower packet loss compared to Europe with the exception of FTTx technology. In actuality, the difference is not significant and is negligible with respect to broadband performance for individual users.

Technology	Europe	US
xDSL	0.46%	0.28%
FTTx	0.16%	0.18%
Cable	0.20%	0.12%

Figure EU.3-42: Comparison of Packet Loss between Europe and the USA, by technology

F Comparison Between Countries - Weighted

F.1 Key Performance Indicators

F.2 Download and Upload Speeds

F.2.1 Download

The metric most commonly associated with broadband performance is download throughput, and it is also the metric that ISPs typically use to advertise their products. Because of this, it receives a large amount of attention from regulators and ISPs when marketing their products.

In order to compare between countries and technologies, which can often have very different performance characteristics, results for download speed are presented as a percentage of advertised speed. This is done primarily to allow the reader to determine more easily how accurate marketing claims of ISPs in certain countries are.

Figures EU.4-1 to EU.4-3 below represent download speed as a percentage of advertised speed for each country considered in this study for xDSL, FTTx and cable technologies. Each figure also shows the average achieved across all countries included in this study. For all technologies, it is typical for most countries to achieve a level of throughput close to the average with few outliers, particularly with regards to xDSL technology. Cable and FTTx perform very well across all countries that meet a minimum sample size, achieving 86.42% and 83.93% of advertised speed on average respectively. Almost all countries using either technology generally perform close to the overall average. Malta is seen to perform the furthest below the average compared to all other countries, although it still remains relatively close to it with 74.48%. This contrasts with the previous measurement period when Malta exceeded its advertised speed.

xDSL technology exhibits a wider distribution of results. Slovakia is again shown to achieve the highest level of throughput as a percentage of advertised speed, with the UK proving to have the lowest performance with 45.01%, far below the average of 71.21%. However, as will be shown later in this study, this has more to do with the advertised rate as opposed to actual performance of broadband in these countries. Most countries' throughput performance as a percentage of advertised speed is closer to the average for FTTx and cable technologies.

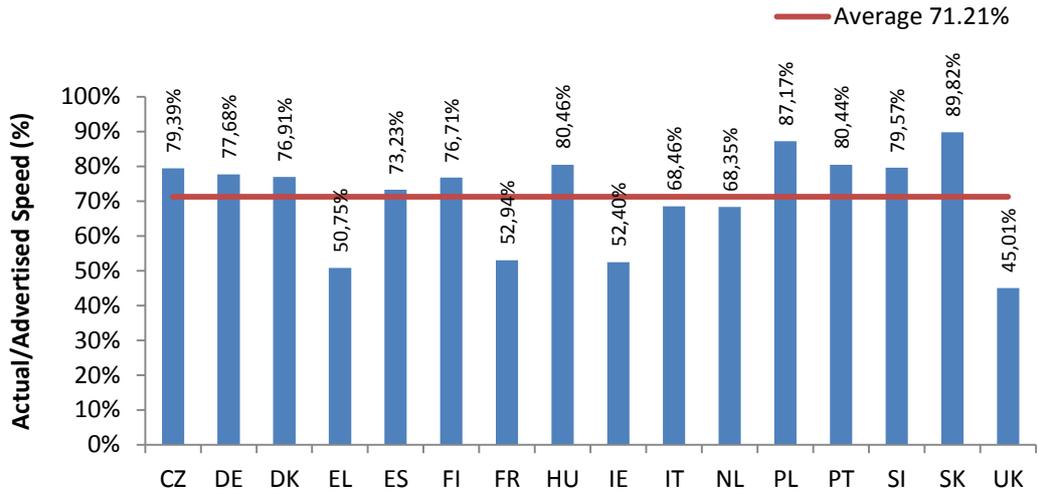


Figure EU.4-1: Weighted Actual Download Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

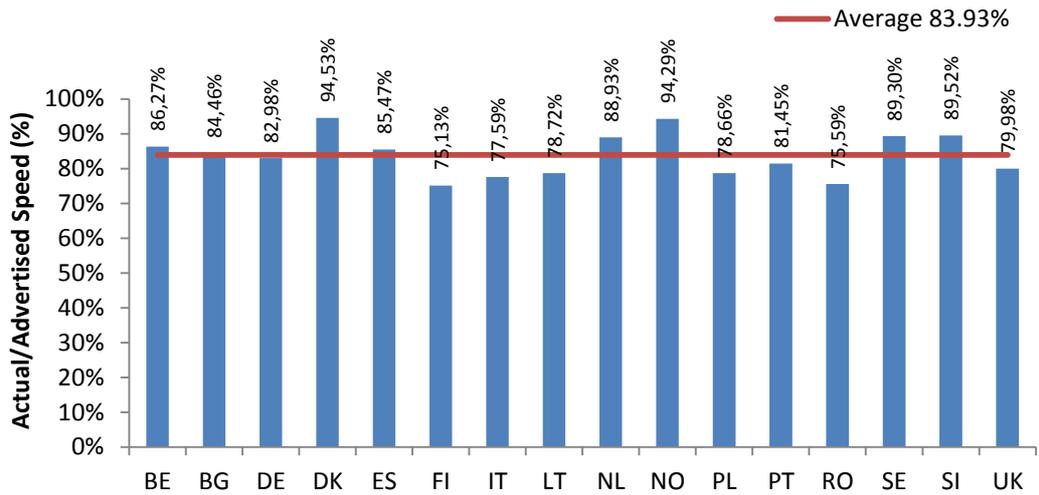


Figure EU.4-2: Weighted Actual Download Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

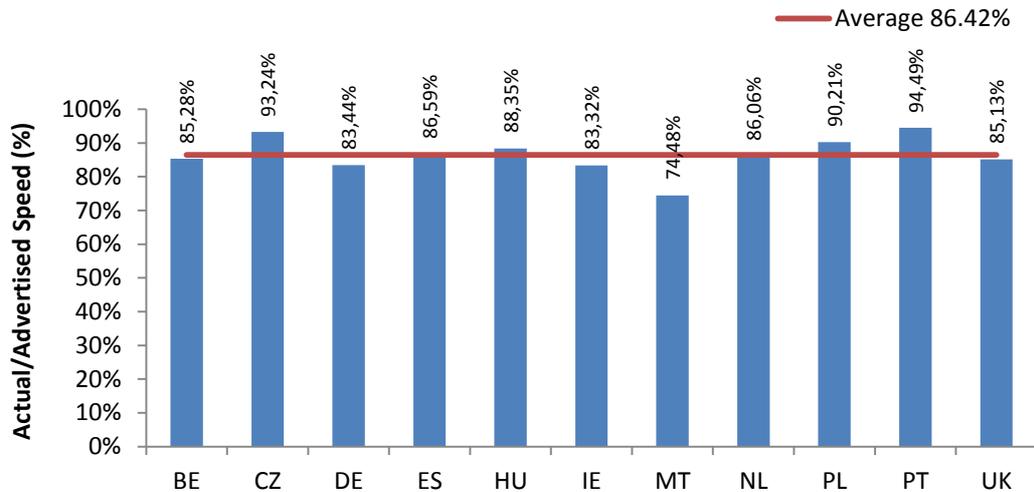


Figure EU.4-3: Weighted Actual Download Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

As mentioned above, nations where cable services are more common tend to achieve figures closer to the advertised rate. While xDSL is quite common throughout Europe, it is much more likely to deliver a lower level of performance compared to cable and FTTx. This is due to access speed degrading with increasing copper loop lengths. Countries such as the UK and France, which deliver some of the poorest results for xDSL, perform better for other technologies, with the UK also exceeding the average and outperforming all countries apart from Ireland for cable.

Overall averages for all access technologies experience an improvement from the previous measurement period of October 2013, particularly cable which shows a 33.9% increase from 47.84Mbps to 64.06Mbps.

Figures EU.4-4 to EU.4-6 below show actual throughput achieved in each country for all access technologies considered in this study. Only countries and technologies with a statistically representative sample are included. The spread of results for actual speed is significantly wider compared to throughput expressed as a percentage of advertised speed, suggesting that similarities in results seen in figures EU.4-1 to EU.4-3 are due to differing marketing strategies in each country. This is showcased particularly by a comparison of actual throughput between the UK and Slovakia for xDSL, who achieve the lowest and highest level of throughput as a percentage of advertised speed respectively although the UK noticeably outperforms Slovakia in real terms. Slovakia also displays one of the lowest levels of actual throughput across all countries.

In the UK and France, xDSL products are predominantly advertised with a single headline speed. Customers use copper phone lines, meaning they can only receive a fraction of the speed advertised by the package. Other countries will offer a wider array of packages and may adopt policies prohibiting providers from selling products customers cannot achieve full speed on.

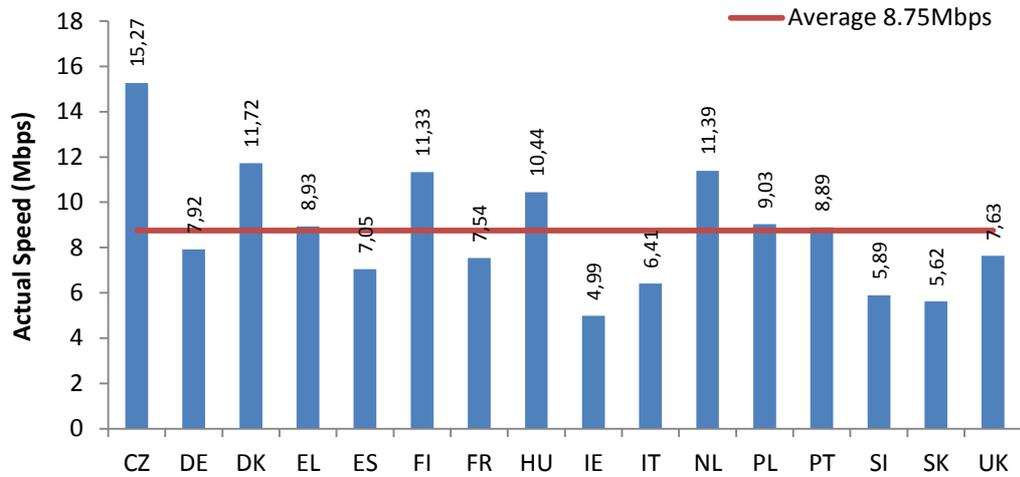


Figure EU.4-4: Weighted Actual Download Speed of xDSL technology during peak periods, by country



Figure EU.4-5: Weighted Actual Download Speed of FTTx technology during peak periods, by country

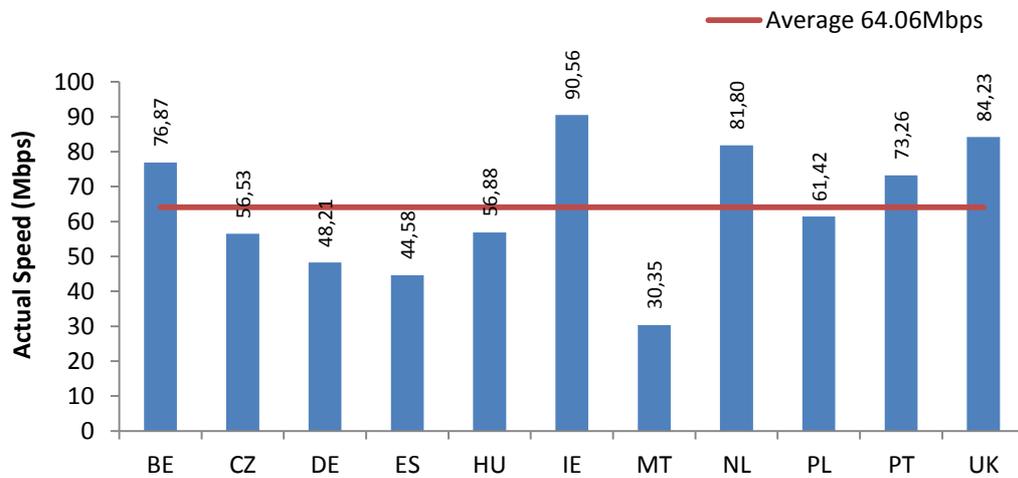


Figure EU.4-6: Weighted Actual Download Speed of cable technology during peak periods, by country

F.2.2 Upload

A great majority of European ISPs provide an asymmetric broadband service, with download throughput usually being many times higher than upload throughput. Historically, this has made sense to ISPs who observed users consumed (downloaded) much more content than they produced and shared (uploaded). However, upload speed is gradually becoming a more important metric as more and more users upload photos, videos and use online storage services. Many ISPs today offer higher upload speed services, recognizing this growing trend.

As with download speed, upload results are presented as a percentage of advertised speed first in order to provide a better comparison between different countries.

Figures EU.4-7 to EU.4-9 display upload speed expressed as a percentage of advertised speed for all access technologies. Only countries and technologies with a statistically representative sample are included. The main thing to note is all technologies achieve on average a greater percentage of advertised upload speed, with the average for cable technology performing just under the advertised speed across all nations. This is likely due to the asymmetry of throughput rates (a service needs to handle less traffic in order to deliver a higher percentage in the upstream direction as the rates are lower). This is particularly important for xDSL services, as the lower upstream target is more manageable even on longer copper phone lines. This may also have to do with lower usage of the upstream direction, although there is not enough data to support this theory.

As mentioned above, cable services perform just below the advertised rate on average, achieving 99.58% of advertised speed. FTTx technology also nearly achieved its advertised rate on average with 95.95%. All access technologies show a slightly lower average than in October 2013, although as with download throughput, this is mostly related to higher advertised rates than actual performance.

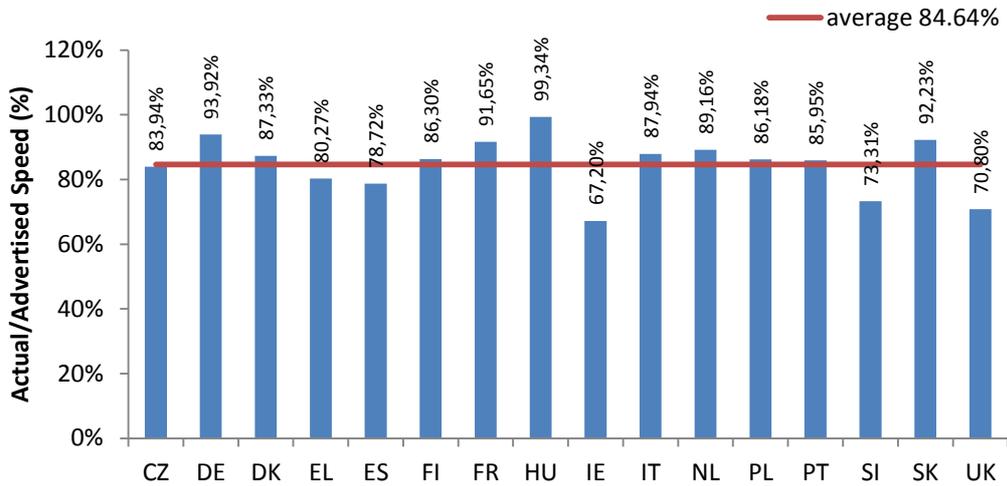


Figure EU.4-7: Weighted Actual Upload Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

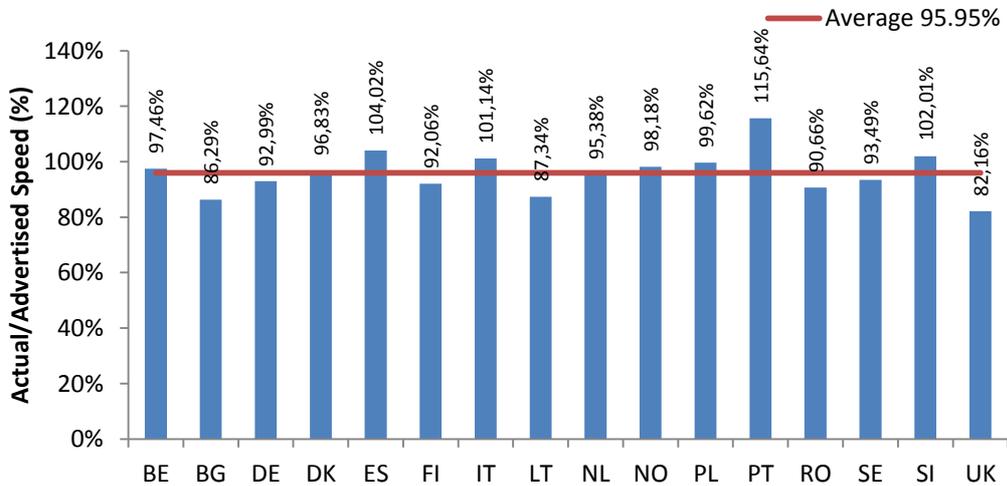


Figure EU.4-8: Weighted Actual Upload Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

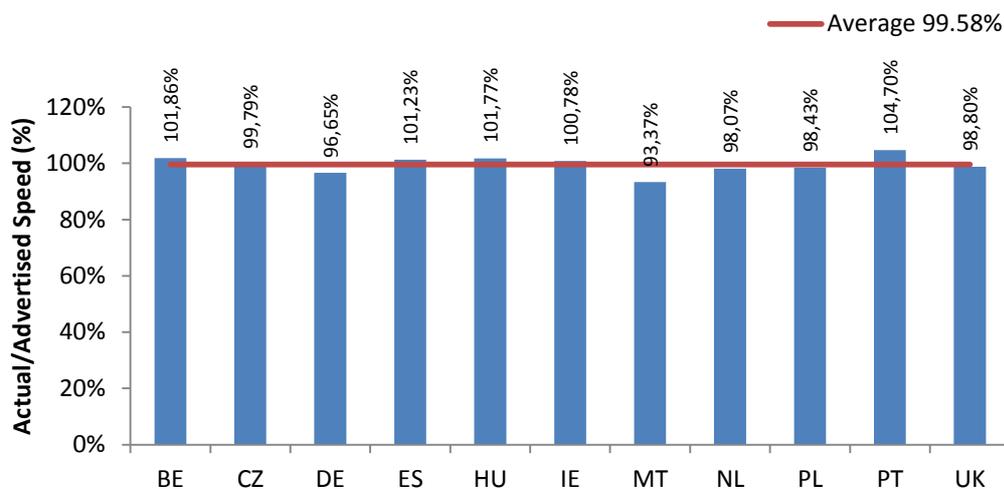


Figure EU.4-9: Weighted Actual Upload Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

Numerous countries exceed their individual advertised rates for upload throughput for both cable and FTTx services, with most others performing very close to the advertised rates. Only Bulgaria, Lithuania and the UK reach below 90% of the advertised speed for FTTx technology. All countries using cable technology exceed 90% of the advertised rates. Numerous countries for xDSL and FTTx technology also exceed the average rate. Portugal especially performed very well across all access technologies, particularly FTTx for which it outperformed all other countries, achieving 115.64% of the advertised speed. It also performs above the average rate of xDSL and cable technologies.

As with download speed, observing actual results for upload speeds by country and access technology shows an improvement in average performance, contrasting with results for throughput expressed as a percentage of advertised speed.

Figures EU.4-10 to EU.4-12 show actual upload speed for each technology across each country. Eastern European and Nordic countries tend to display higher levels of throughput compared to their western counterparts, as was the case in October 2013.

Lithuania outperforms all other countries significantly with regards to FTTx technology in spite of the fact their performance in upload speed as a percentage of advertised speed was one of the lowest compared to other countries.

As was also observed in October 2013, the asymmetric relationship between download and upload speed is most evident in xDSL technology. With an average download speed of 8.75Mbps and upload speed of 0.84Mbps, the ratio between download and upload is approximately 10:1. Cable technology also shows a ratio of nearly 10:1, contrasting with the previous measurement period, whereas FTTx technology continues to display a ratio of roughly 2:1.

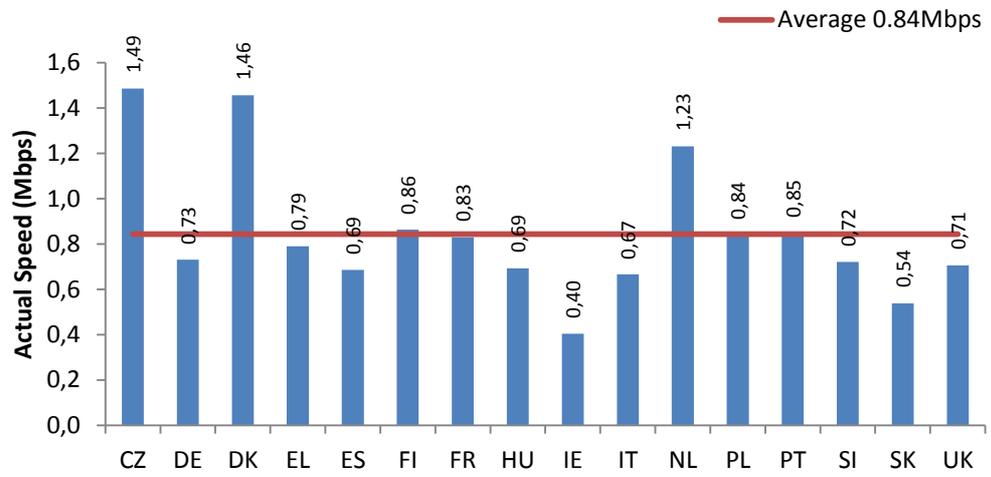


Figure EU.4-10: Weighted Actual Upload Speed of xDSL technology during peak periods, by country

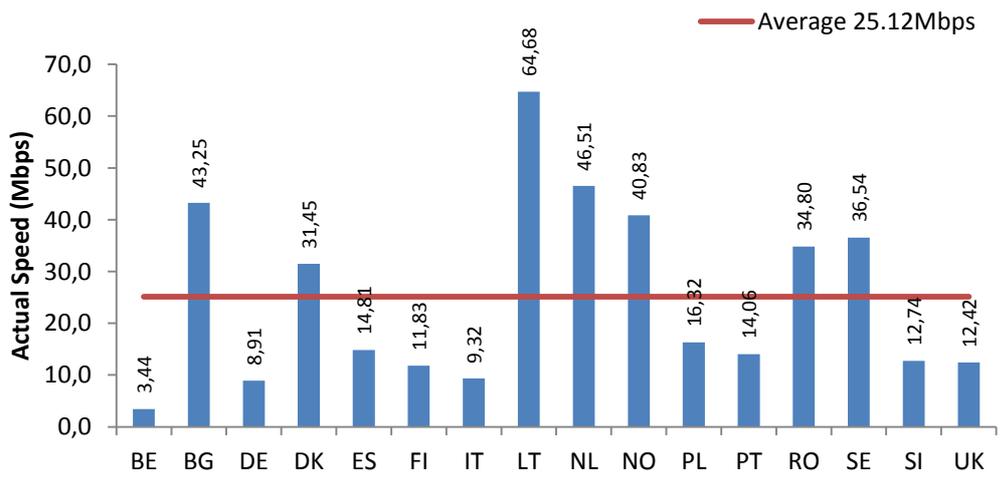


Figure EU.4-11: Weighted Actual Upload Speed of FTTx technology during peak periods, by country

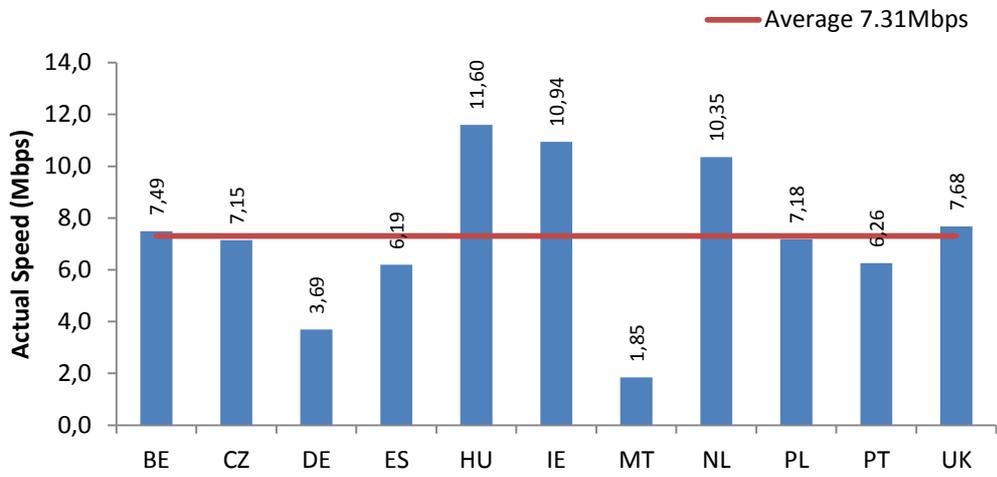


Figure EU.4-12: Weighted Actual Upload Speed of cable technology during peak periods, by country

F.2.3 Latency

Latency is an important metric often overlooked by consumers. Latency, also referred to as round-trip latency, is a measure of how long it takes for a single packet of data to go from point A to point B and back. In this study, round-trip latency is measured between the panelists' homes and the nearest measurement server.

Given that every communication with Internet services involves the transmission and reception of packets of data, latency affects everything we do on the Internet. It is especially important for time-sensitive applications such as online gaming, video streaming and voice communications. The lower the latency, the faster and 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 is most typically the dominant factor affecting levels of latency.

Figures EU.4-13 to EU.4-15 show the average round-trip latency per country, split by each access technology. On average, cable delivered the lowest latency of 20.49ms, compared to FTTx's 22.11ms and xDSL's 36.36ms. FTTx technology exhibits a slightly increased average latency since October 2013, with xDSL and cable showing a small improvement.

Thanks to the deployment of FTTH technology in certain eastern European regions, countries such as Bulgaria and Slovenia delivered the best latency performances for FTTx technology. It should also be noted that latency of these two countries also outperforms the latency of all countries using cable technology despite cable achieving a lower overall average. This is because FTTH technology does not need to use an xDSL based last mile technology that causes a significant latency overhead.

Of particular note is the significantly higher latency of xDSL technology from Spain, which also delivered the highest latency results in October 2013. This is despite the fact there are measurement servers located in Spain (Madrid).

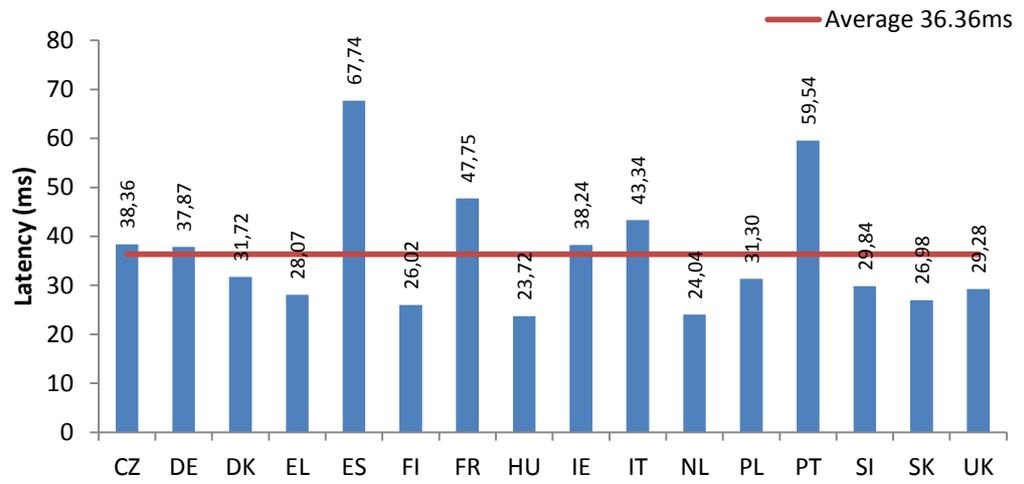


Figure EU.4-13: Weighted Latency of xDSL technology during peak periods, by country

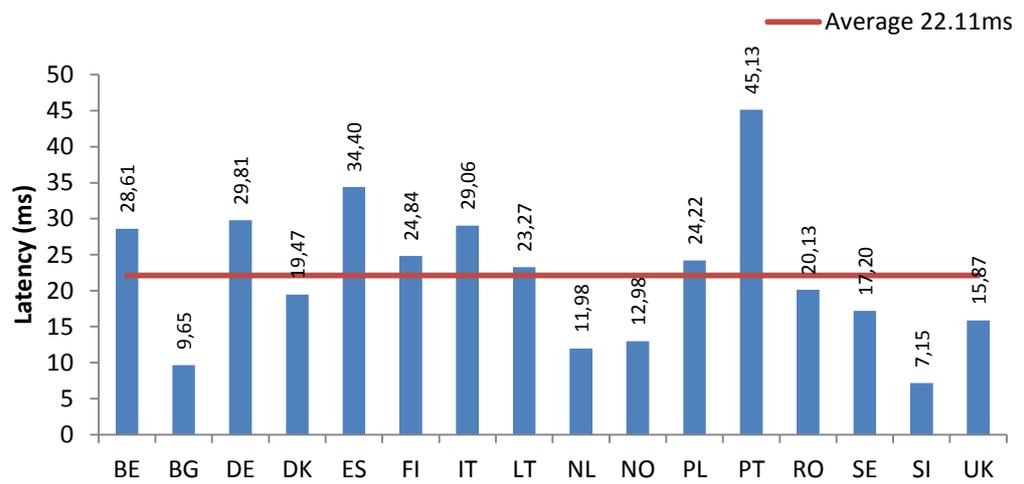


Figure EU.4-14: Weighted Latency of FTTx technology during peak periods, by country

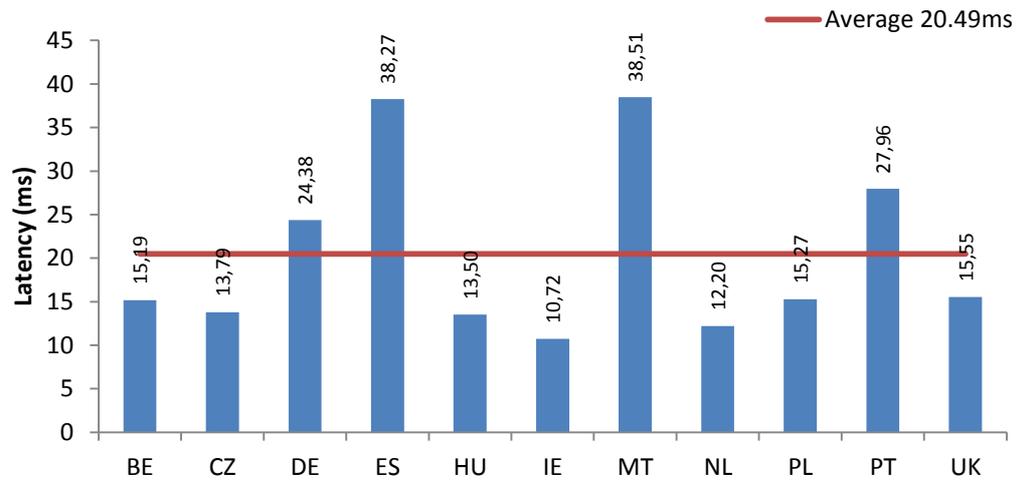


Figure EU.4-15: Weighted Latency of cable technology during peak periods, by country

F.2.4 Packet Loss

As can be the case with latency, packet loss is often overlooked by consumers. This metric describes what percentage of packets sent from the home to the server they are communicating with and back are lost during the transmission. When packets are lost, the two parties involved in the communication will usually attempt to retransmit data in order to account for the loss. This takes time and becomes very noticeable to users when it reaches a certain level.

Realtime applications such as online gaming, video streaming and voice communications are the most affected by high packet loss.

Figures EU.4-16 to EU.4-18 show packet loss figures by country, split for each type of access technology during the peak period. As was the case with packet loss during the previous measurement period, most countries exhibit very low packet loss with only two countries meeting or exceeding 0.75%. Most countries shown to exceed 0.5% packet loss are seen on xDSL technology, notably Greece, Hungary, Ireland, France and Italy. Italy also shows significantly higher packet loss than all other countries on FTTx technology. FTTx and xDSL technologies both show slightly lower averages since October 2013, in contrast to cable technology which displays a slightly higher average packet loss. This is due to a number of countries exhibiting higher packet loss. Examples include Portugal, Poland and Spain. Differences in packet loss between each country are negligible in real world terms.

It is not unusual for packet loss on xDSL technology to be higher than it is for all other access technologies due to the use of older copper lines. These lines are more likely to suffer from physical defects which may inhibit communications and thus overall performance. Otherwise, most nations across all access technologies show packet loss ranging between 0.1% and 0.4% and also do not exceed the overall averages.

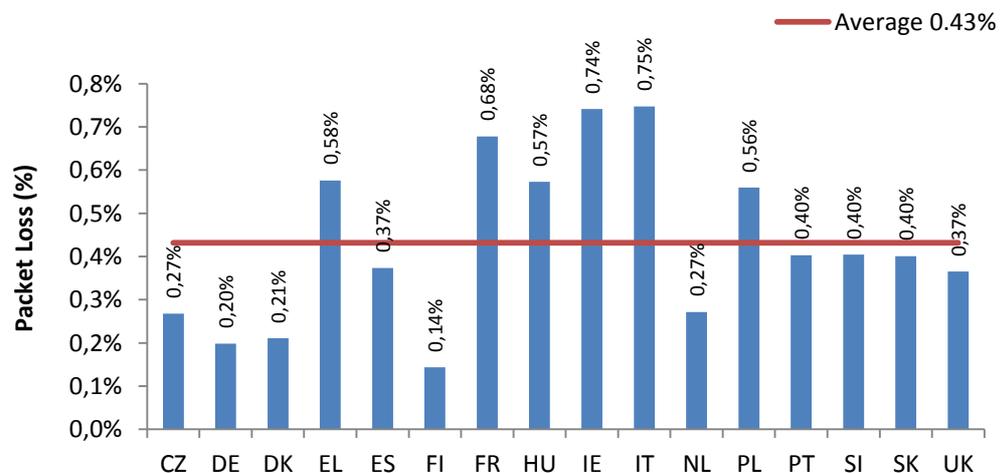


Figure EU.4-16: Weighted Packet loss of xDSL technology during peak periods, by country

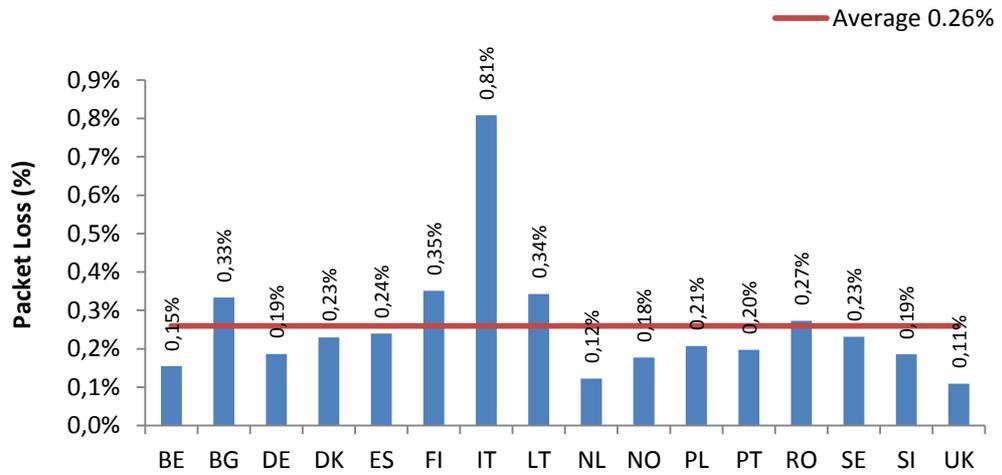


Figure EU.4-17: Weighted Packet loss of FTtx technology during peak periods, split by country

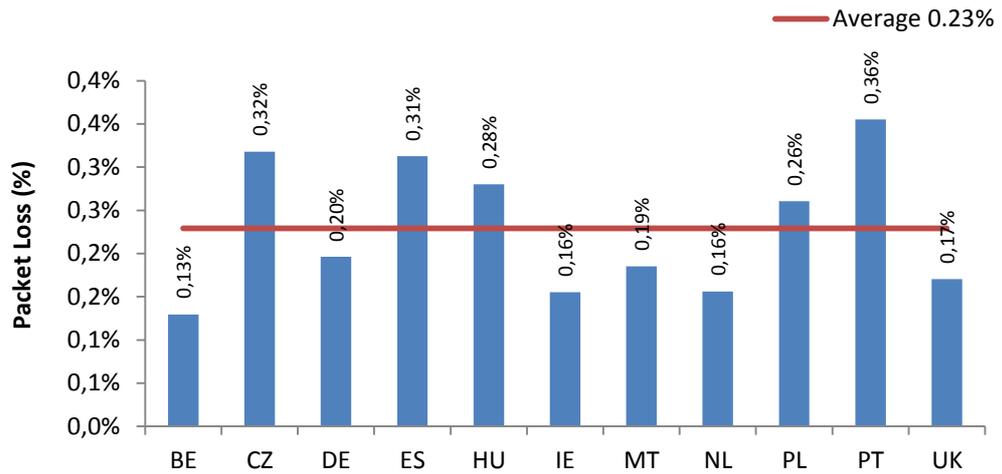


Figure EU.4-18: Weighted Packet loss of cable technology during peak periods, by country

F.2.5 DNS Resolution Time and Failure Rate

DNS is a very important Internet service that allows you to turn hostnames, e.g. `www.youtube.com`, into IP addresses that your computer can communicate with. DNS services are typically provided by the ISP to provide a fast, nearby service for their users to use. A DNS service performing badly can lead users to perceive noticeable delays. This is especially evident during web browsing activities, which rely extensively on DNS.

Theoretically, a good DNS deployment should provide DNS resolution times and failure rates that are at most equal to latency and packet loss figures respectively. This is because DNS servers are typically hosted inside the ISP's networks. Thus, this traffic doesn't need to leave the ISP's network.

Figures EU.4-19 to EU.4-21 show DNS resolution time for all technologies split by country during peak hours. As was the case in October 2013, FTTx technology delivered the best DNS resolution times, averaging 17.69ms, with those nations deploying FTTH technology exhibiting the best resolution times of all. Examples include Bulgaria, Lithuania and Slovenia, all of which also outperform countries using cable technology. The average DNS resolution time for cable technology is almost the same at 18.47ms, with xDSL exhibiting a significantly higher average resolution time of 35.37ms.

As was the case in the previous measurement period, Belgium proves to be an exception to other countries using FTTH technology, with its FTTx (VDSL in this case) services averaging much higher DNS resolution times than latency. Other nations whose DNS resolution times don't closely resemble their latency results include Hungary, Italy and Ireland for xDSL technology and Poland for xDSL and FTTx technologies, with DNS resolution times proving to be much higher.

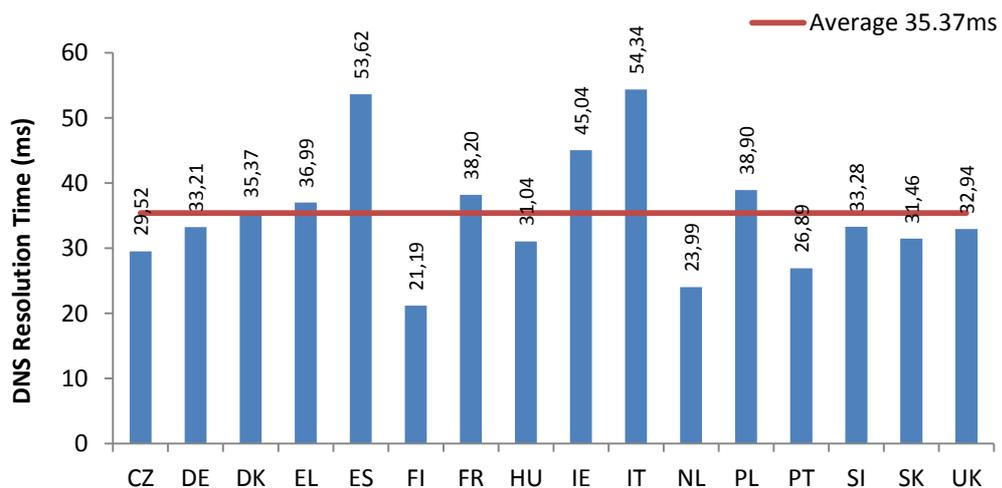


Figure EU.4-19: Weighted DNS Resolution Time of xDSL technology during peak periods, by country

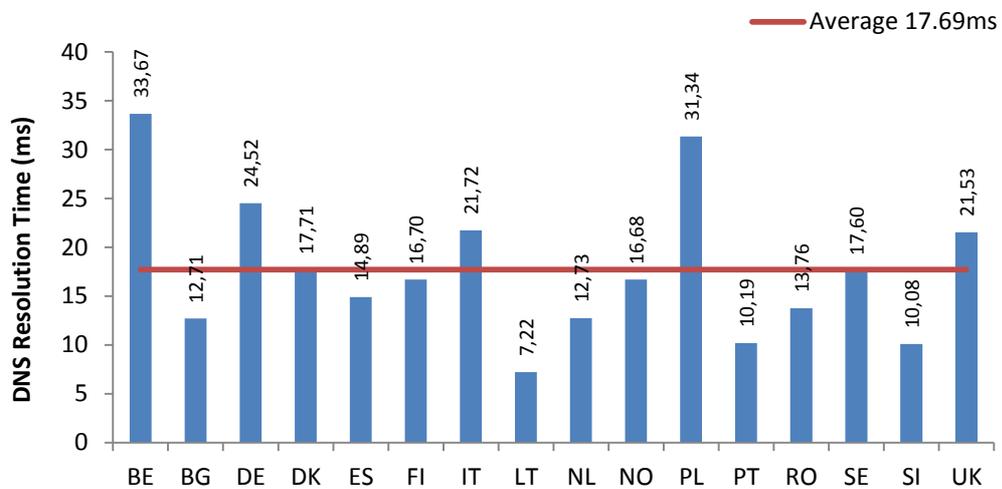


Figure EU.4-20: Weighted DNS Resolution Time of FTTx technology during peak periods, by country

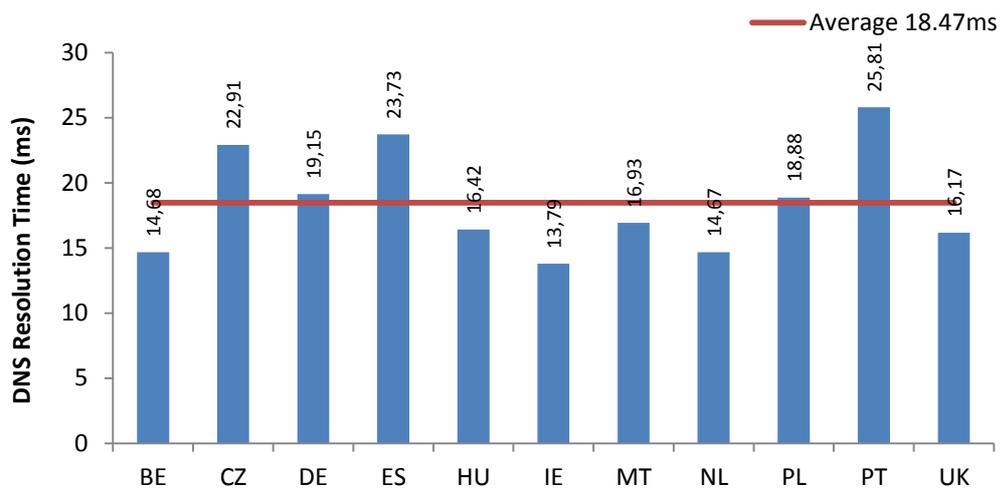


Figure EU.4-21: Weighted DNS Resolution Time of cable technology during peak periods, by country

Figures EU.4-22 to EU.4-24 show DNS failure rate per country for each access technology during the peak period. On average, xDSL and cable technologies display lower failure rates than in October 2013, with most countries displaying lower DNS failure rates as well. Exceptions include Slovenia for xDSL technology and Ireland for cable technology.

In contrast, DNS failure rate on FTTx technology is slightly higher on average. Some countries, particularly Bulgaria, Norway and Finland, displayed a significant increase in failure rates, with most other countries reporting a slight increase. Exceptions include Spain and Sweden, which demonstrate a noticeable improvement compared to the previous year.

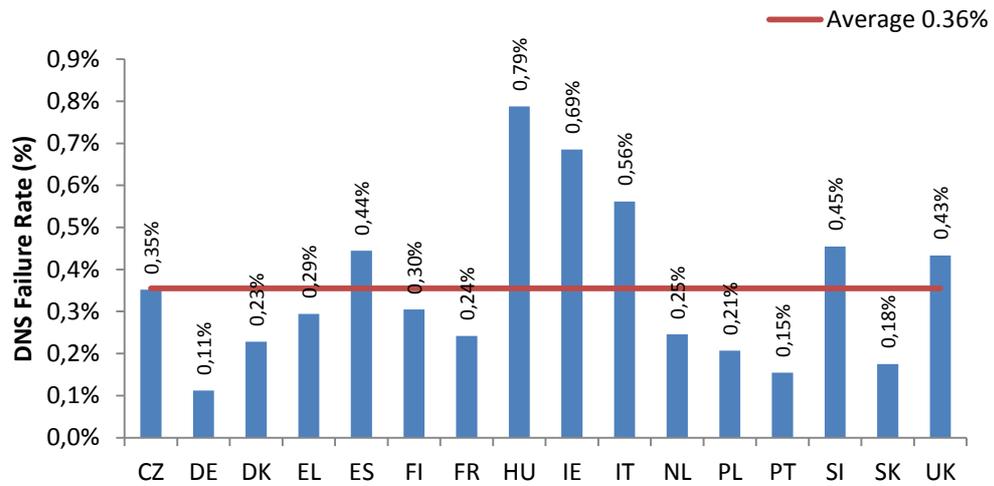


Figure EU.4-22: Weighted DNS Resolution Failure Rate of xDSL technology during peak periods, by country

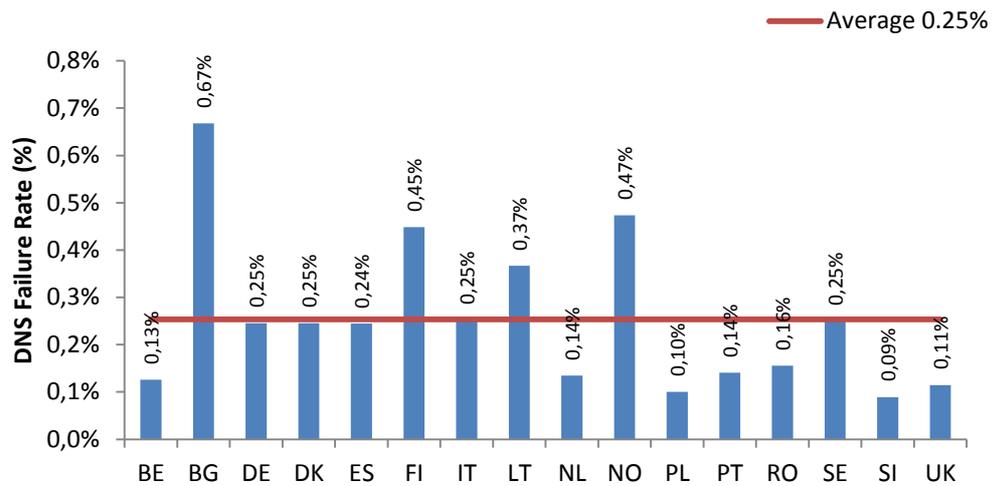


Figure EU.4-23: Weighted DNS Resolution Failure Rate of FTTx technology during peak periods, by country

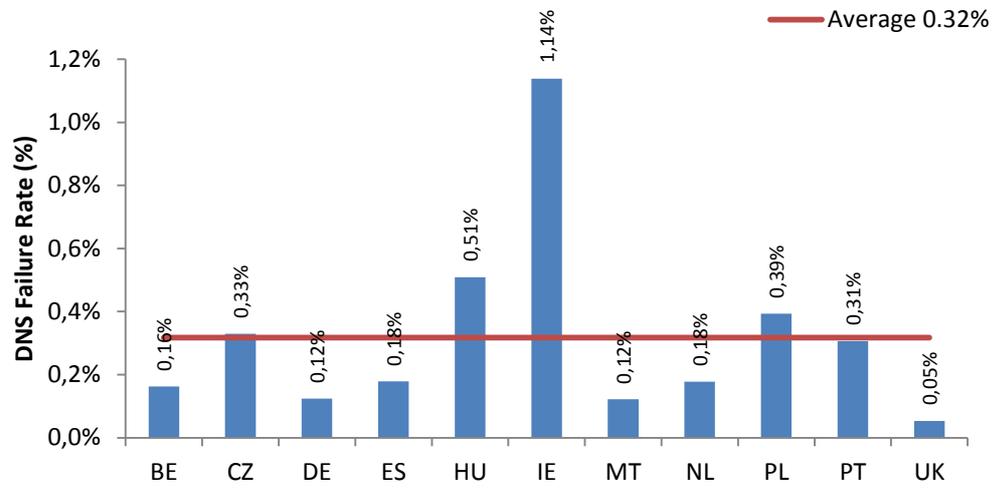


Figure EU.4-24: Weighted DNS Resolution Failure Rate of cable technology during peak periods, by country

F.2.6 Web Browsing Speeds

Figures EU.4-25 to EU.4-27 below display webpage loading times in each country during the peak period, split by access technologies. The test was performed to the public-facing websites of Facebook, Google and YouTube, whose servers are geographically hosted across Europe to optimize consumer performance.

Webpage loading time is between 1 and 3 seconds for most countries using xDSL technology with Ireland as the sole exception (as was the case during the previous testing period), exhibiting a loading time of 3.94 seconds. Loading times are consistently below 1 second for countries making use of Cable and FTTx technologies. Some exceptions include Lithuania and Slovenia for FTTx technology as well as Spain and Malta for cable technology. Poland's loading times also exceed 1 second for both FTTx and cable technologies.

Countries that exhibited excellent actual download throughput also display low webpage loading times, a fact made clear with the lower average loading times for cable and FTTx technologies. However, web browsing speed does not improve proportionately with download throughput. This is because webpage loading time is not just a function of line speed, but also latency, and for services offering 10Mbps download speed or more, latency dominates web browsing performance. This can be seen through the average webpage loading times of cable and FTTx technologies being almost identical. The average loading time for cable based services is slightly higher than FTTx as well, although the difference is negligible. All access technologies have shown a small increase in their average performances since October 2013.

Some exceptions to the above exist where countries with relatively high latency exhibit lower loading times compared to other countries. For example, Portugal, which demonstrated a very high latency for FTTx services also reports one of the lowest loading times. The opposite may also occur, with countries showing lower latency also displaying higher loading times, such as FTTx services in Norway and Bulgaria. This implies the connectivity between our test websites and the Portuguese ISPs is better than to our measurement servers, and the opposite is true in the case of Norway and Bulgaria.

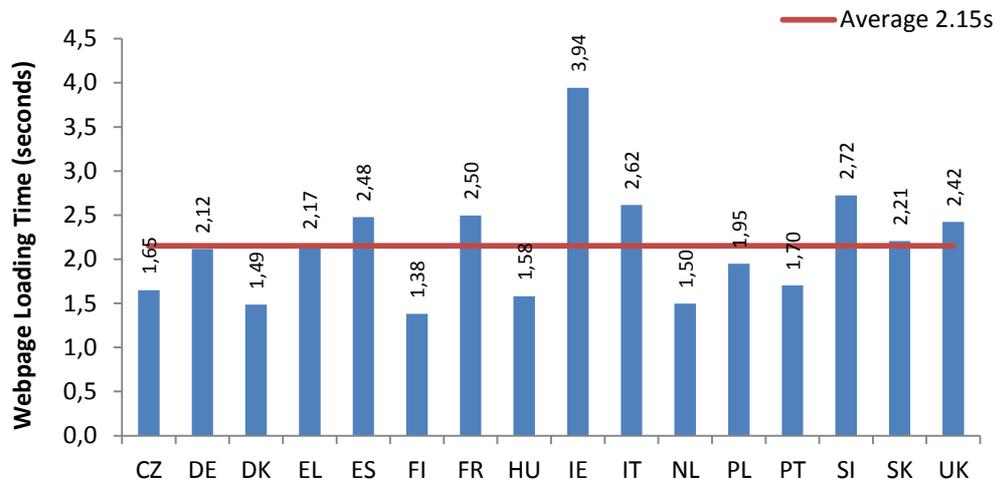


Figure EU.4-25: Weighted Webpage Loading Time of xDSL technology during peak periods, by country

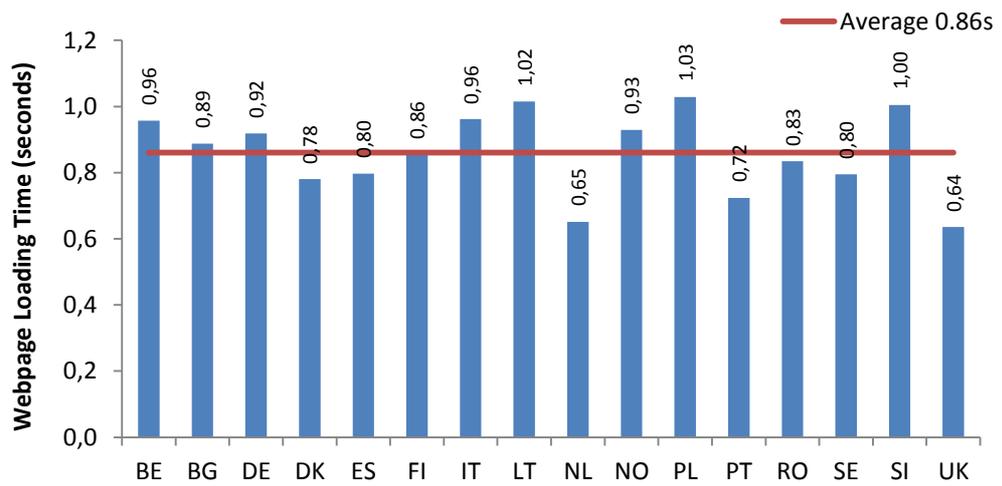


Figure EU.4-26: Weighted Webpage Loading Time of FTTx technology during peak periods, by country

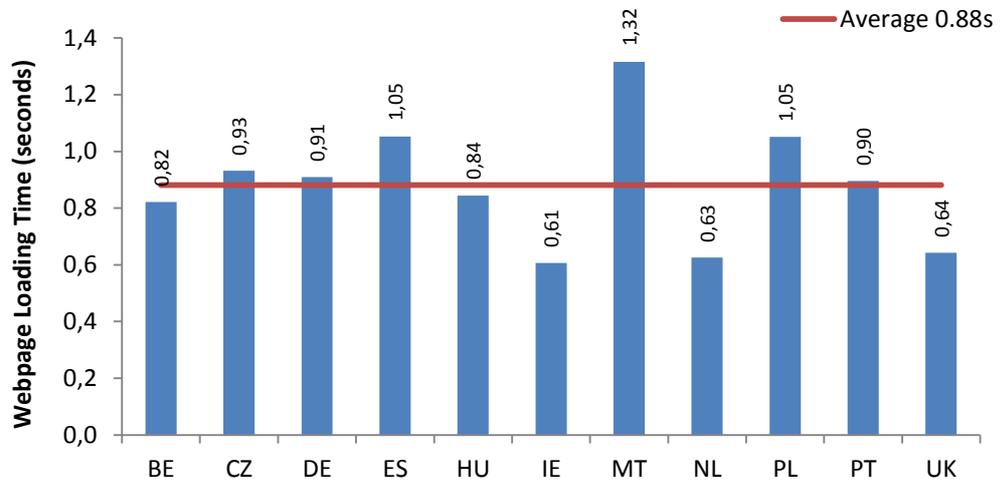


Figure EU.4-27: Weighted Webpage Loading Time of cable technology during peak periods, by country

F.2.7 VoIP Jitter

Jitter is an important metric for users who frequently use realtime communication applications. It can also be referred to as latency consistency. Broadband connections frequently shifting between 10ms and 20ms latency would have a high jitter value. This pattern would be very noticeable to consumers using realtime applications such as video streaming and online gaming. Thus, it is better to have lower jitter.

This study reports on downstream and upstream jitter separately. Both are important for two-way communications such as phone calls, but significant technological differences make it so results in the downstream and upstream directions are noticeably divergent.

Downstream jitter is shown in figures EU.4-28 to EU.4-30 during the peak period for each access technology, split by country. Jitter in the downstream direction is very low for most countries across all access technologies. Only xDSL services in a handful of countries show an average significantly above 1ms, such as Ireland, Germany, France, Portugal, Italy and Slovenia, with xDSL technology also displaying the highest average compared to cable and FTTx based services. The downstream jitter performance of xDSL services improved greatly in Poland and Slovakia since the previous testing period. Average downstream jitter is slightly higher for FTTx and xDSL based services, in contrast to cable technology which shows a lower average since October 2013. The lower level of jitter for cable technology compared to FTTx technology also reflects the behaviour of latency on each service.

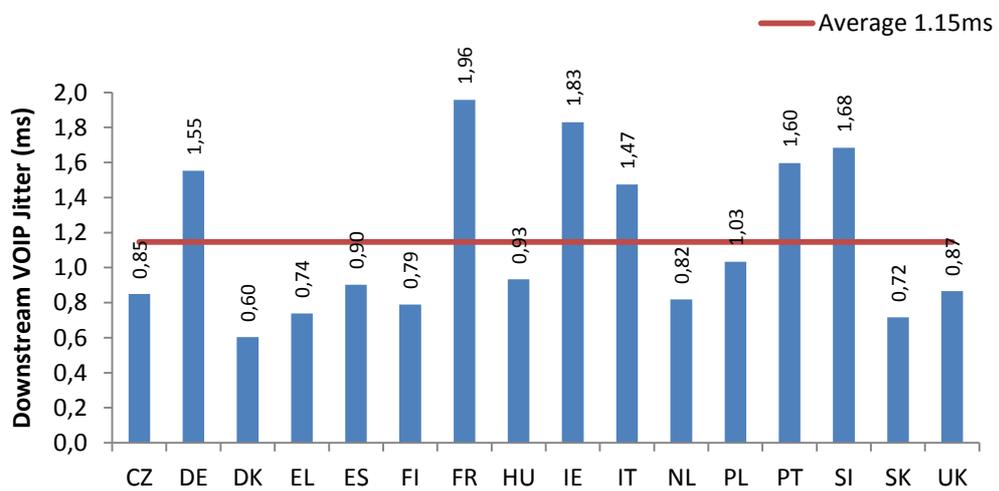


Figure EU.4-28: Weighted Downstream VoIP Jitter of xDSL technology during peak periods, by country

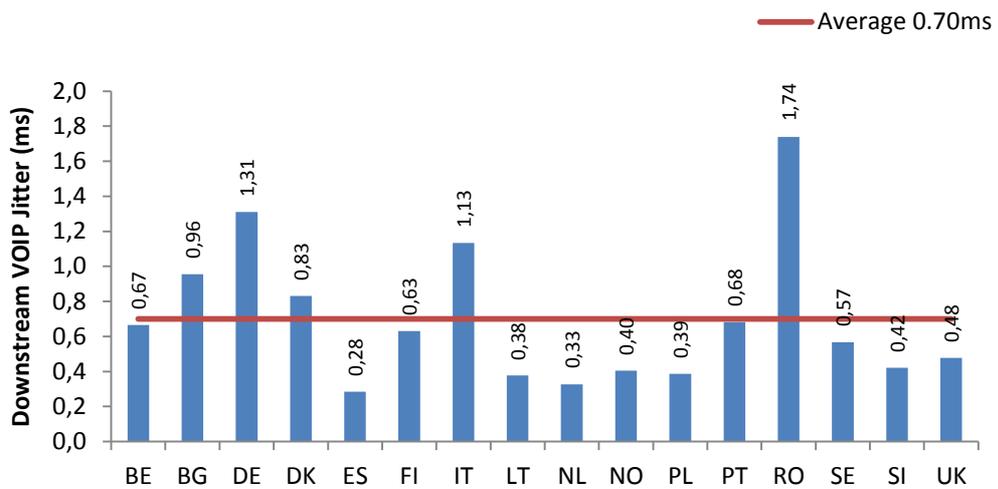


Figure EU.4-29: Weighted Downstream VoIP Jitter of FTTx technology during peak periods, by country

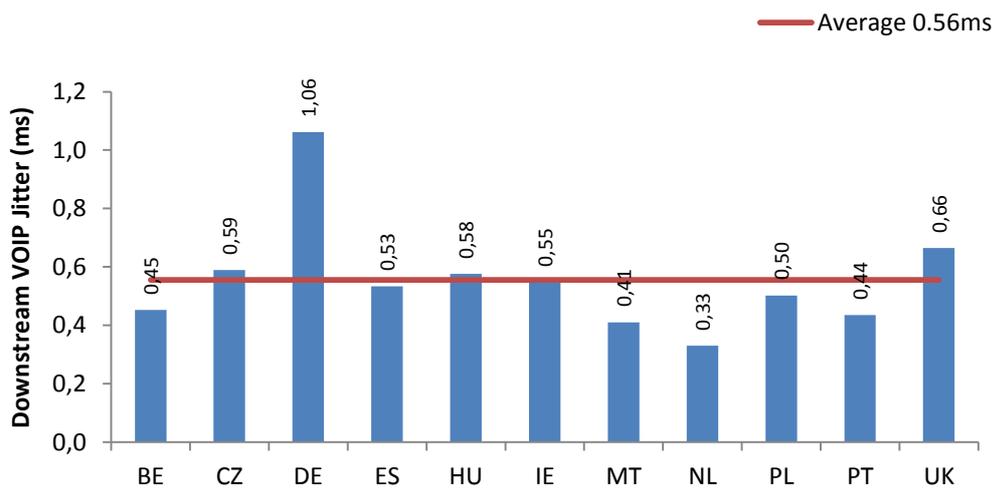


Figure EU.4-30: Weighted Downstream VoIP Jitter of cable technology during peak periods, by country

Figures EU.4-31 to EU.4-33 show results for upstream jitter during the peak period for all access technologies, split by country. Most nations using xDSL services also exhibited similar jitter results, with only a few exceeding the overall average. Germany's upstream jitter for xDSL services also improved greatly, but most other nations display higher upstream jitter than in October 2013. This contrasts with FTTx services as upstream jitter rarely exceeded 2ms, with few countries showing upstream jitter higher than 1ms.

Average upstream jitter of cable technology also improved since October 2013, although it still displays the highest average upstream jitter of 3.02ms. Initially, this may appear to be a contradiction with downstream jitter results as well as most other metric performances. The reason for cable's high upstream jitter is due to the fact that they are based on the concept of TDMA (Time Division Multiple Access). This means the modem's time is divided in slots during which it can either send or receive data, but not both simultaneously. If the modem is busy while the user tries to send a packet, the packet will have to wait in a queue until an opportunity to be transmitted arises. This can result in small but common variations in packet delays.

It is also important to remember that although upstream jitter is relatively high for cable networks, its actual level is low enough that it can be deemed negligible with regards to overall broadband performance. For instance, most Voice over IP (VoIP) phones have a de-jitter buffer of at least 25ms. This means jitter under 25ms would not affect the call quality. As can be seen in figures below, the upstream jitter cable technology across each country is far below 25ms.

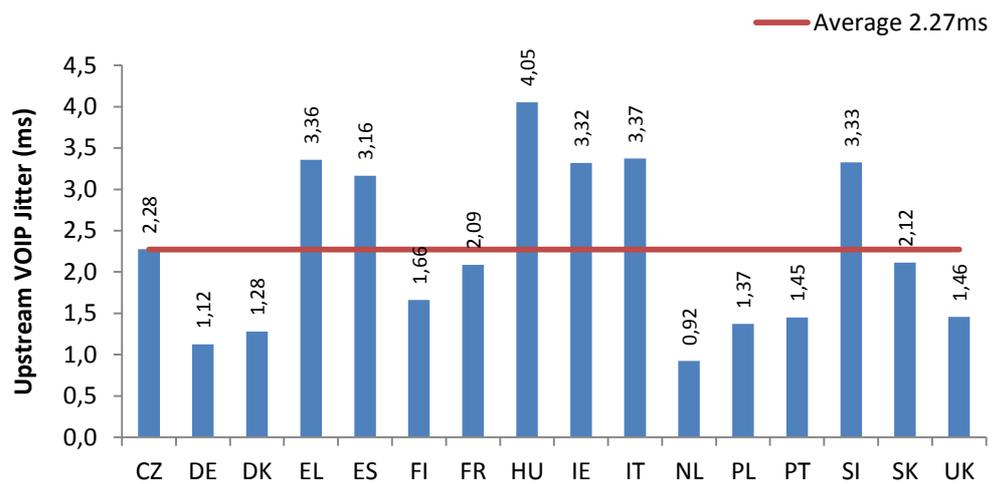


Figure EU.4-31: Weighted Upstream VoIP Jitter of xDSL technology during peak periods, split by country

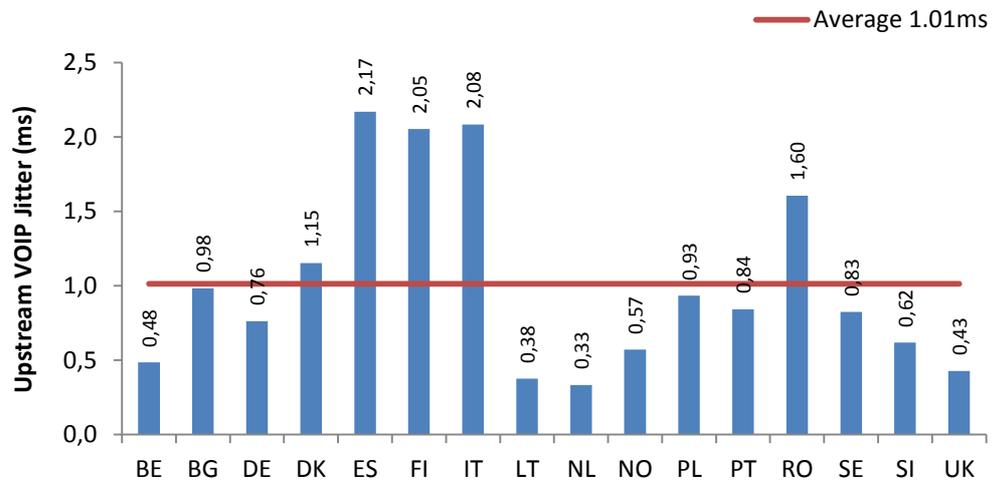


Figure EU.4-32: Weighted Upstream VoIP Jitter of FTTx technology during peak periods, by country

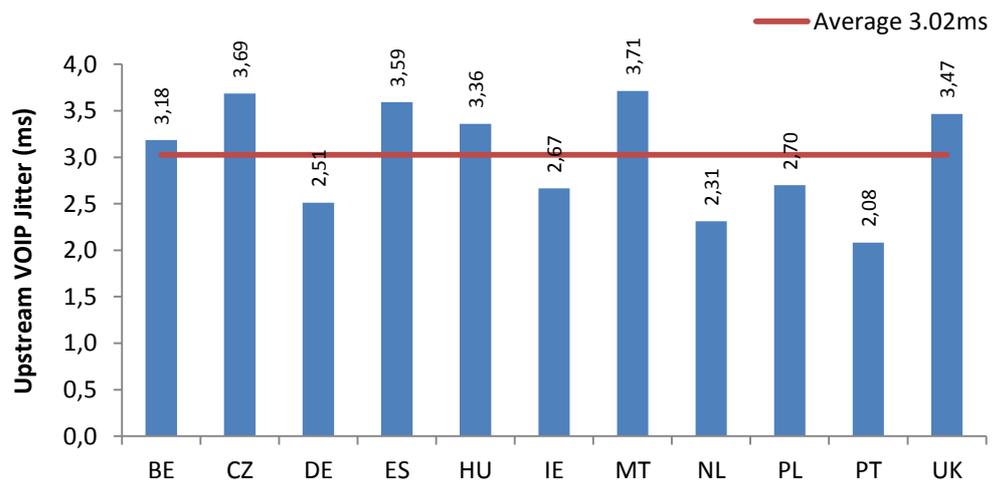


Figure EU.4-33: Weighted Upstream VoIP Jitter of cable technology during peak periods, by country

Section G Overview

In Section G, data is weighted by technology, ISP and speed tier market share. This market share data was supplied by member state's NRAs. The granularity varied by NRA, and only the following countries supplied data in sufficient detail to enable inclusion in this section: Bulgaria, Czech Republic, Hungary, Malta, the Netherlands, and the UK.

Weights are calculated by taking the number of subscribers on a specific speed tier, on an ISP with a given technology in an individual country. Each panellist's measurements are multiplied by this figure, and the resulting country level average is obtained by summing these values and then dividing through the sum of the weights. This prevents bias toward other technologies that may or may not perform better than the one in question.

In some cases the NRAs did not supply market share figures for a speed tier and technology combination which was included in our sample. It is difficult to isolate the cause of all such cases, but a likely cause could be that the ISP excluded this tier from the data they provided, perhaps because it represented too small a fraction of their user base or it was a legacy product.

Such cases affected the weightings for 6.31% of panellists included in section G. In these cases we applied a neutral weight to these panellists' data. Neutral weights are obtained by calculating the average market share of each technology across all ISPs and speed tiers. The small percentage affected leads us to state that the results presented in section G are still representative for the countries included.

Please note that this approach differs from Section F, which weighs only by ISP subscriber count. The differences in the results presented in Section F and Section G demonstrate that the sample used was significantly over-represented in higher-speed tiers in some countries. For example, in Section G xDSL results in the Czech Republic are 43% lower, FTTx results are 24% lower in Bulgaria, and cable results are 38% lower in Malta and 17% lower in the UK.

Throughput in the downstream and upstream directions is lower for xDSL technology in the Czech Republic, FTTx technology in Bulgaria and cable technology in the UK once the above approach is employed. This is because lower speed tiers make up a larger market share in these countries. Throughput of cable technology in the UK is also lower as its weighted average is no longer influenced by other ISP's market shares.

G Comparison Between Countries - Weighted

G.1 Key Performance Indicators

G.2 Download and Upload Speeds

G.2.1 Download

Figures EU.5-1 to EU.5-3 below represent download speed as a percentage of advertised speed for each country and technology. For all technologies, it is typical for most countries to exceed 80% of the headline speeds apart from xDSL technology, although most countries still exceed 75% of advertised rates. Malta is seen to perform below all other countries for cable technology, achieving 79.44% of the advertised speed, similarly to its performance depicted in sections D and F. This also contrasts with the previous measurement period when Malta exceeded its advertised speed.

xDSL technology exhibits a wider distribution of results. Slovakia is again shown to achieve the highest level of throughput as a percentage of advertised speed, with the UK proving to have the lowest performance with 45.75. However, as will be shown later in this study, this has more to do with the advertised rate as opposed to actual performance of broadband in these countries. Countries exhibit more similar results compared to each other for cable and FTTx technologies.

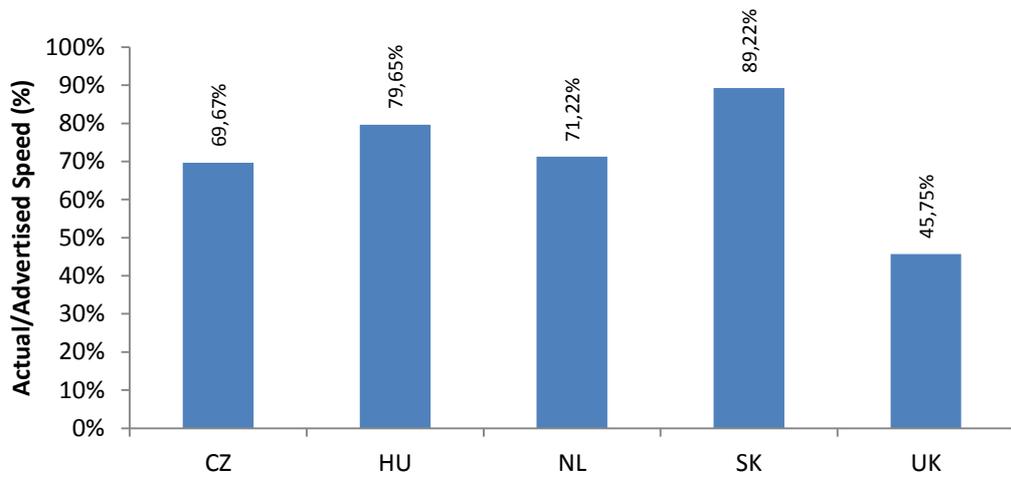


Figure EU.5-1: Weighted Actual Download Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

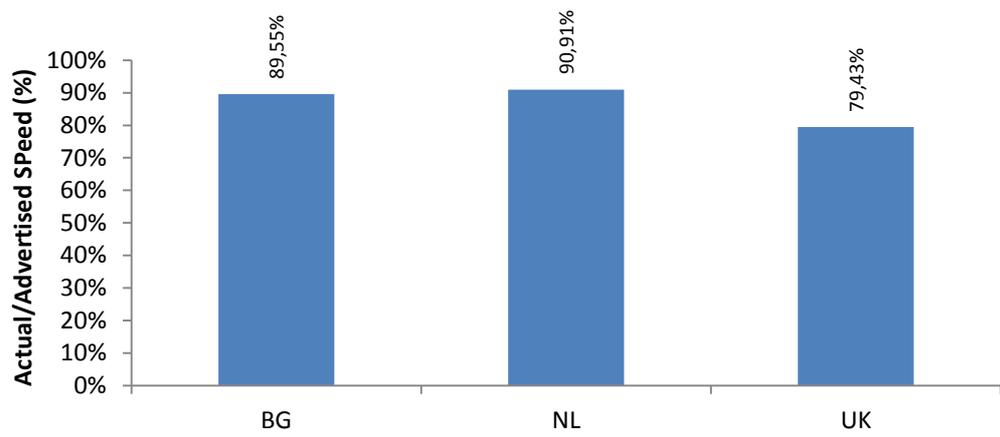


Figure EU.5-2: Weighted Actual Download Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

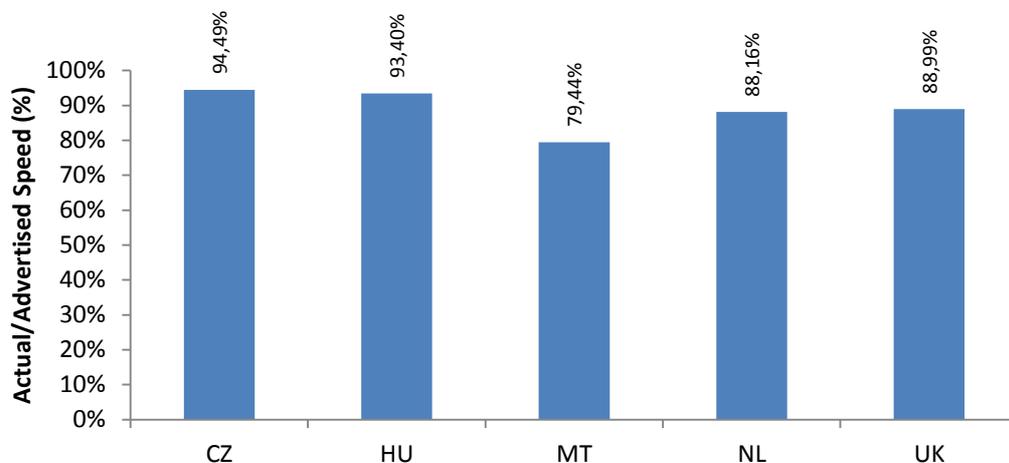


Figure EU.5-3: Weighted Actual Download Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

As mentioned above, nations where cable services are more common generally tend to achieve figures closer to the advertised rate. While xDSL is quite common throughout Europe, it is much more likely to deliver a lower level of performance compared to cable and FTTx. This is due to access speed degrading with increasing copper loop lengths. Countries such as the UK, which deliver some of the poorest results for xDSL, perform better for other technologies.

Figures EU.5-4 to EU.5-6 below show actual throughput achieved in each country for all access technologies considered in this study. The spread of results for actual speed is significantly wider compared to throughput expressed as a percentage of advertised speed, particularly for cable technology, suggesting that similarities in results seen in figures EU.5-1 to EU.5-3 are due to differing marketing strategies in each country. This is showcased particularly by a comparison of actual throughput between the UK and Slovakia for xDSL, who achieve the lowest and highest level of throughput as a percentage of advertised speed respectively although the UK noticeably outperforms Slovakia in real terms. Slovakia also displays the lowest levels of actual throughput across all countries.

In the UK, xDSL products are predominantly advertised with a single headline speed. Customers use copper phone lines, meaning they can only receive a fraction of the speed advertised by the package. Other countries will offer a wider array of packages and may adopt policies prohibiting providers from selling products customers cannot achieve full speed on.

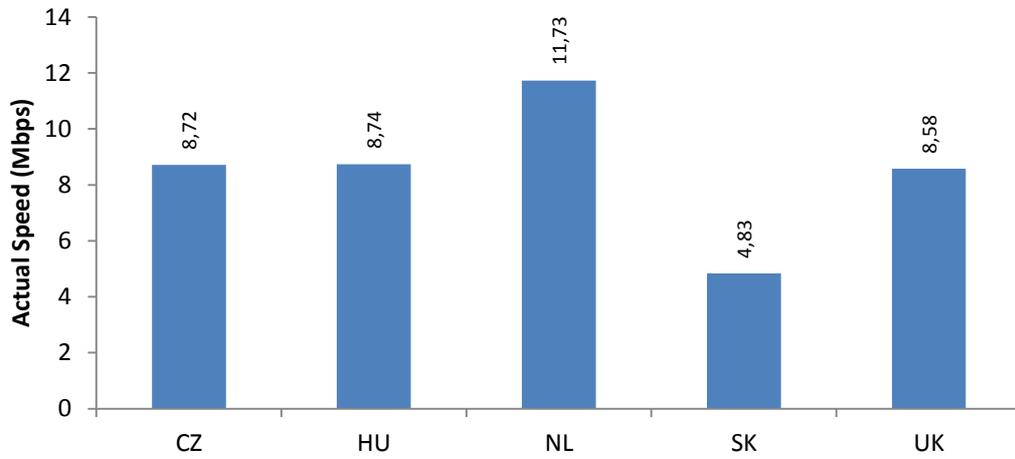


Figure EU.5-4: Weighted Actual Download Speed of xDSL technology during peak periods, by country

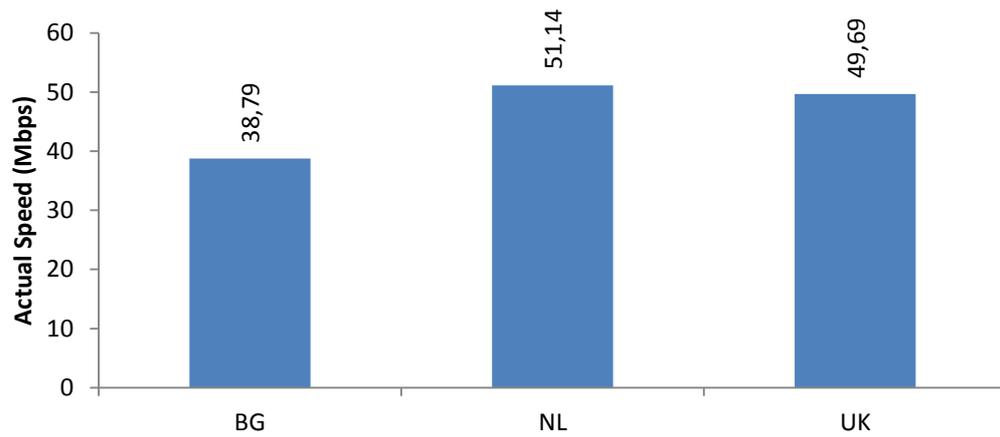


Figure EU.5-5: Weighted Actual Download Speed of FTTx technology during peak periods, by country

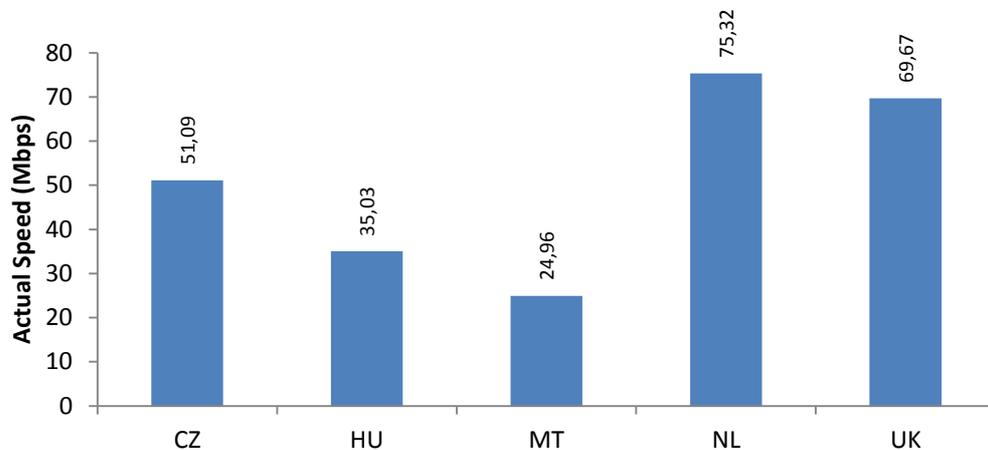


Figure EU.5-6: Weighted Actual Download Speed of cable technology during peak periods, by country

6.2.2 Upload

Figures EU.4-7 to EU.4-9 display upload speed expressed as a percentage of advertised speed for all access technologies. Only countries and technologies with a statistically representative sample are included. The main thing to note is all countries achieve a greater percentage of advertised upload speed across all technologies, with xDSL technology in the Czech Republic being the sole exception. This is likely due to the asymmetry of throughput rates (a service needs to handle less traffic in order to deliver a higher percentage in the upstream direction as the rates are lower). This is particularly important for xDSL services, as the lower upstream target is more manageable even on longer copper phone lines. This may also have to do with lower usage of the upstream direction, although there is not enough data to support this theory.

Most countries perform very close to their respective headline speeds across all technologies, with some also exceeding headline speeds. Hungary does so across both cable and xDSL technologies, with cable technology in the Czech Republic also exceeding the advertised speed. Upload throughput of cable technology is generally higher across all countries compared to xDSL and FTTx. Very few countries across all technologies achieve below 90% of the advertised speed, mostly for xDSL technology. Malta is the only country using cable technology which performs just short of 90% of advertised speed.

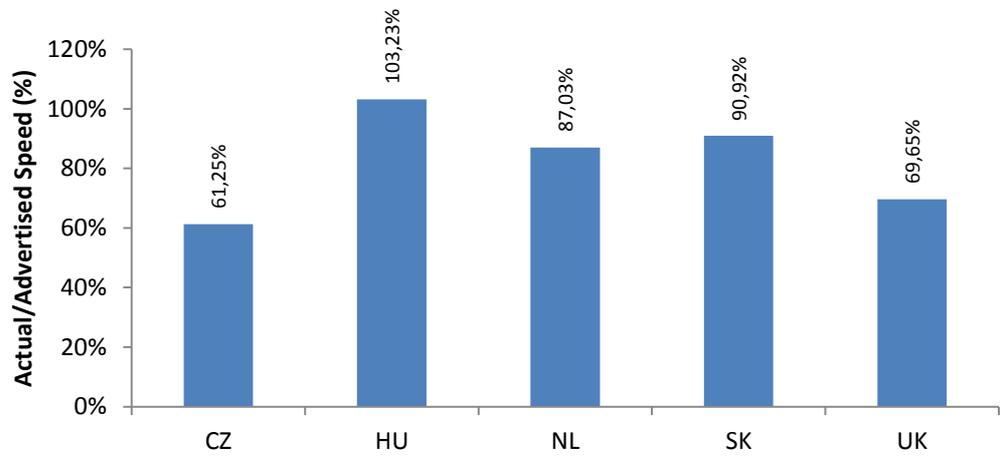


Figure EU.5-7: Weighted Actual Upload Speed of xDSL technology as a Percentage of Advertised Speed during peak periods, by country

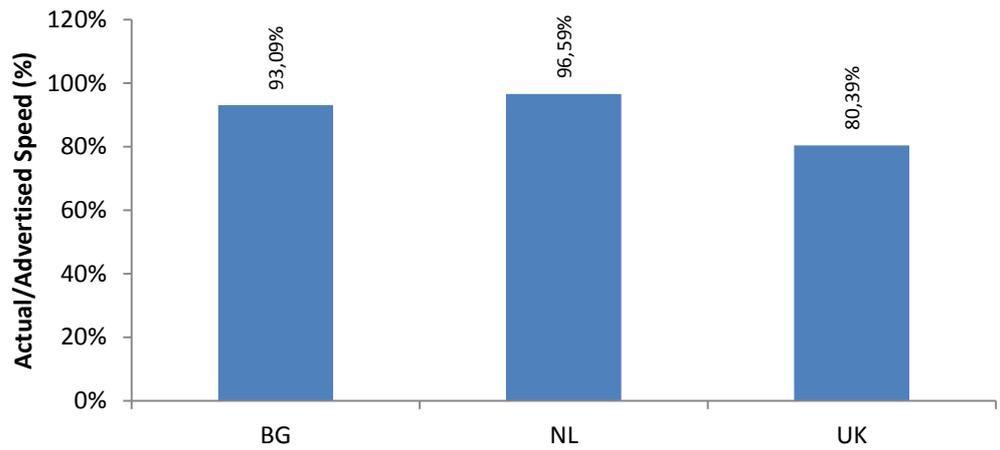


Figure EU.5-8: Weighted Actual Upload Speed of FTTx technology as a Percentage of Advertised Speed during peak periods, by country

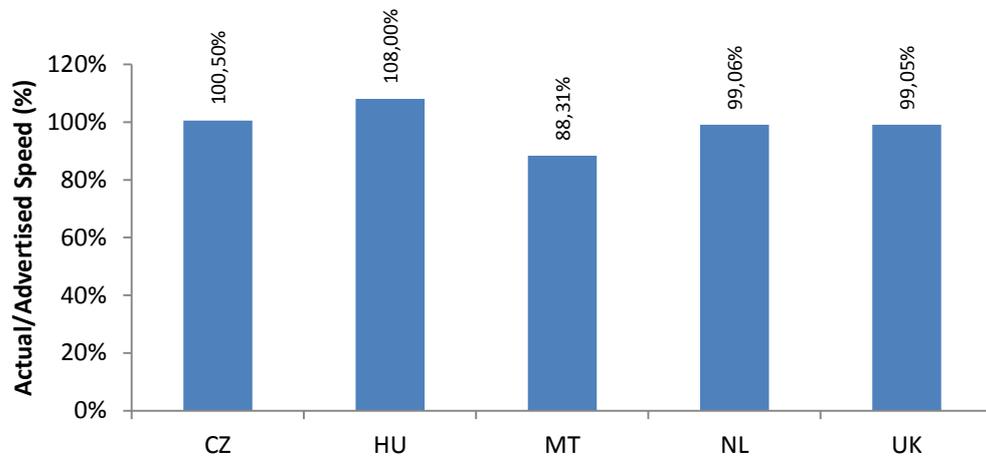


Figure EU.5-9: Weighted Actual Upload Speed of cable technology as a Percentage of Advertised Speed during peak periods, by country

Figures EU.5-10 to EU.5-12 show actual upload speed for each technology across each country. Eastern European and Nordic countries tend to display higher levels of throughput compared to their western counterparts.

The Netherlands outperforms all other countries across all technologies, contrasting with upload speed of xDSL technology in Hungary which proves to be one of the lowest in spite of outperforming all other countries in percentage terms. Like with download speed, actual upload speed across each country also displays a larger spread than upload speed expressed as a percentage of advertised speed. This also suggests similarities in percentage results are due to a diverse marketing strategy across each nation.

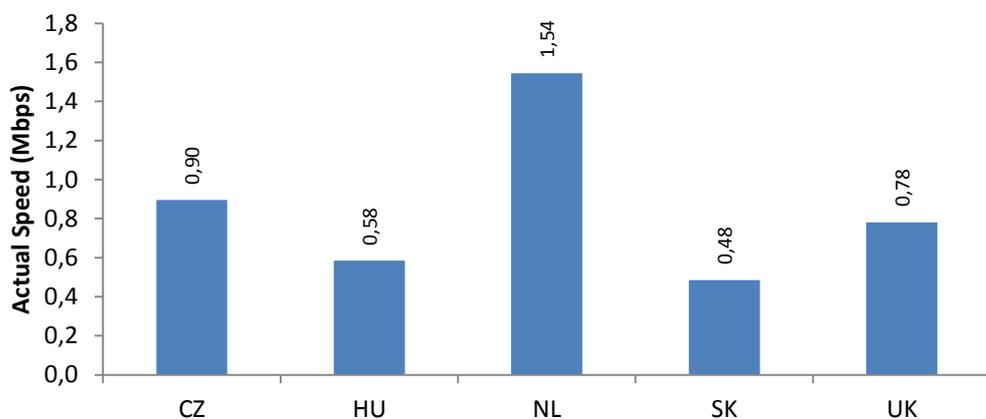


Figure EU.5-10: Weighted Actual Upload Speed of xDSL technology during peak periods, by country

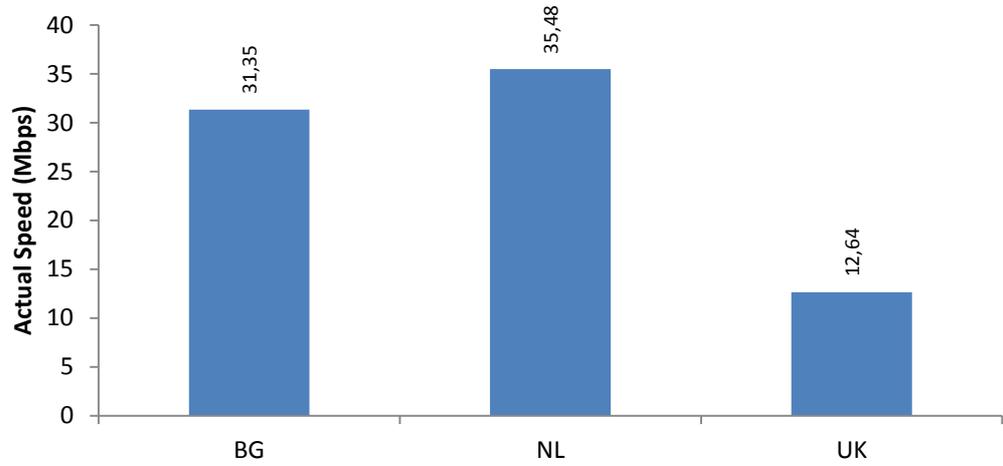


Figure EU.5-11: Weighted Actual Upload Speed of FTTx technology during peak periods, by country

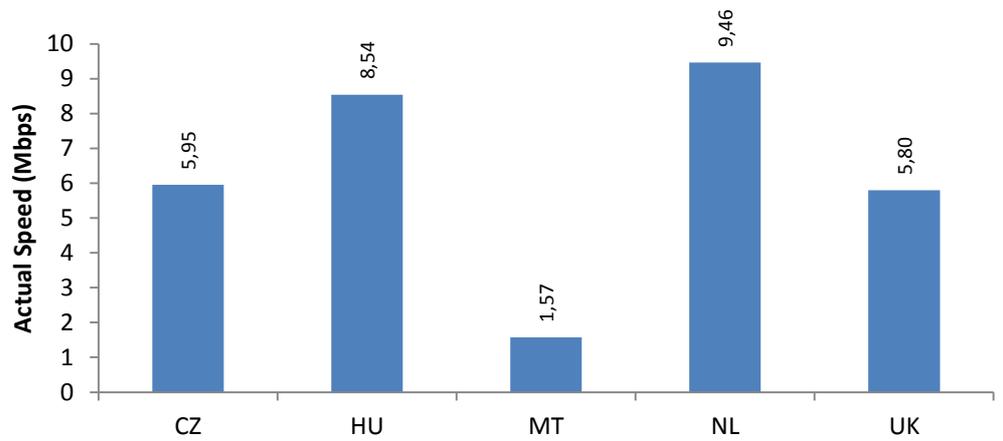


Figure EU.5-12: Weighted Actual Upload Speed of cable technology during peak periods, by country

G.2.3 Latency

Figures EU.5-13 to EU.5-15 show the average round-trip latency per country, split by each access technology. In general, latency proves lowest across most countries providing cable technology, with the exception of Malta which displays a very high latency of 38.93ms. This reflects the behaviour of throughput in this country, which proved lower than all other nations providing cable technology. Latency of xDSL technology across each country is generally higher than latency of all other technologies.

As mentioned previously in this study, the deployment of FTTH technology in certain eastern European regions, countries in said regions delivered the best latency for FTTx technology. Similarly to what is shown in sections D and F, Bulgaria delivered the best latency across all countries providing FTTx technology and also outperforms latency of countries providing cable technology. This is because FTTH technology does not need to use an xDSL based last mile technology that causes a significant latency overhead.

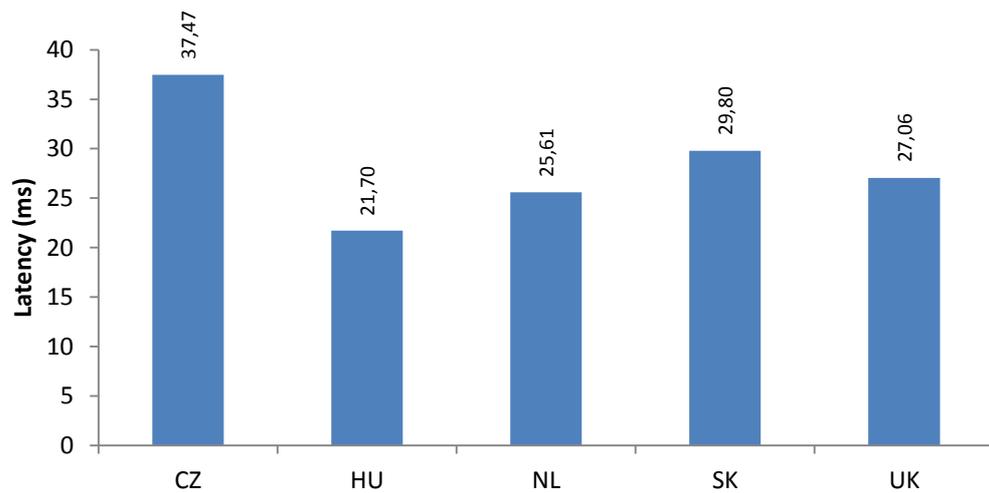


Figure EU.5-13: Weighted Latency of xDSL technology during peak periods, by country

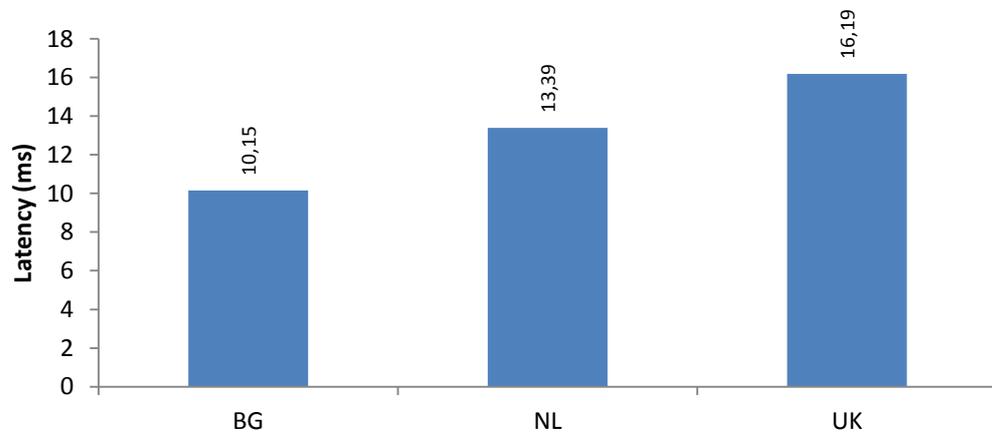


Figure EU.5-14: Weighted Latency of FTTx technology during peak periods, by country

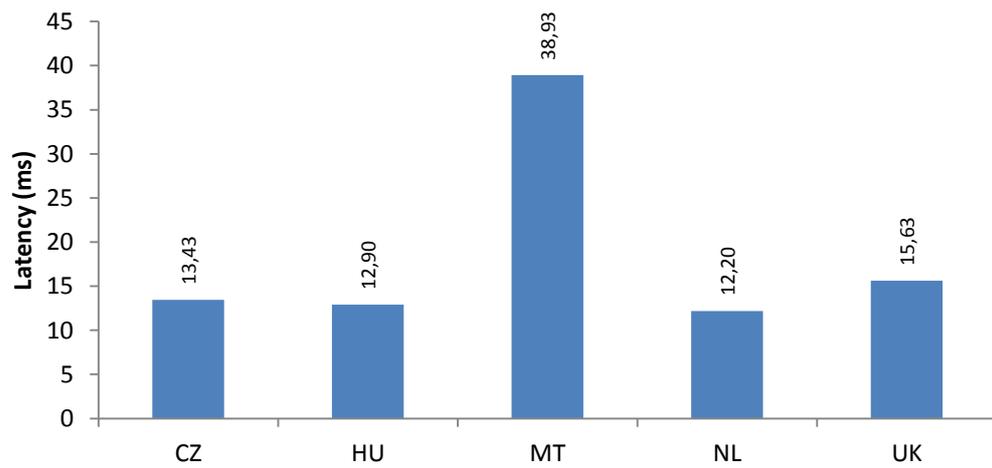


Figure EU.5-15: Weighted Latency of cable technology during peak periods, by country

G.2.4 Packet Loss

Figures EU.5-16 to EU.5-18 show packet loss figures by country, split for each type of access technology during the peak period. Most countries exhibit very low packet loss with only a handful of countries exceeding 0.3%. Most countries shown to exceed 0.3% packet loss are seen on xDSL technology, including the Czech Republic, Slovakia and the UK. Bulgaria also shows a relatively higher packet loss compared to the Netherlands and the UK for FTTx technology.

It is not unusual for packet loss on xDSL technology to be higher than it is for all other access technologies due to the use of older copper lines. These lines are more likely to suffer from physical defects which may inhibit communications and thus overall performance.

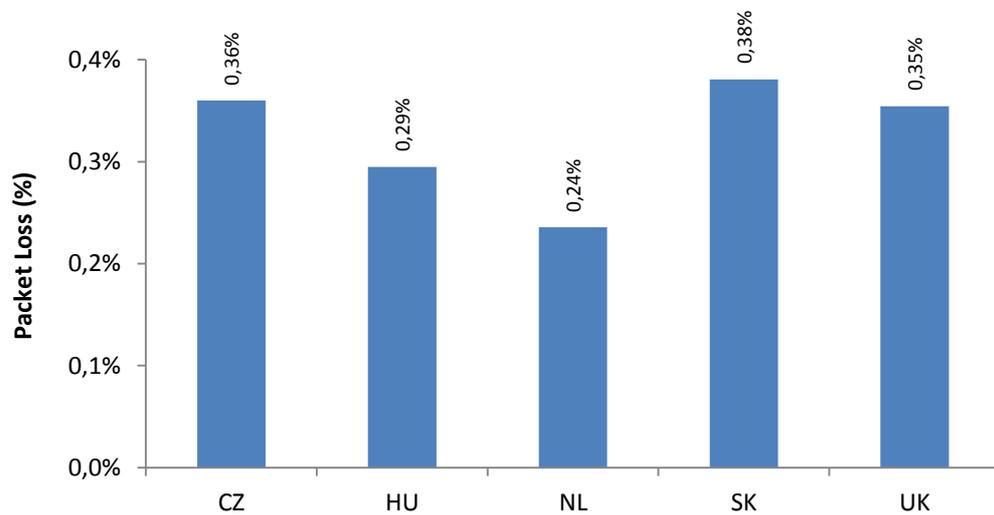


Figure EU.5-16: Weighted Packet loss of xDSL technology during peak periods, by country

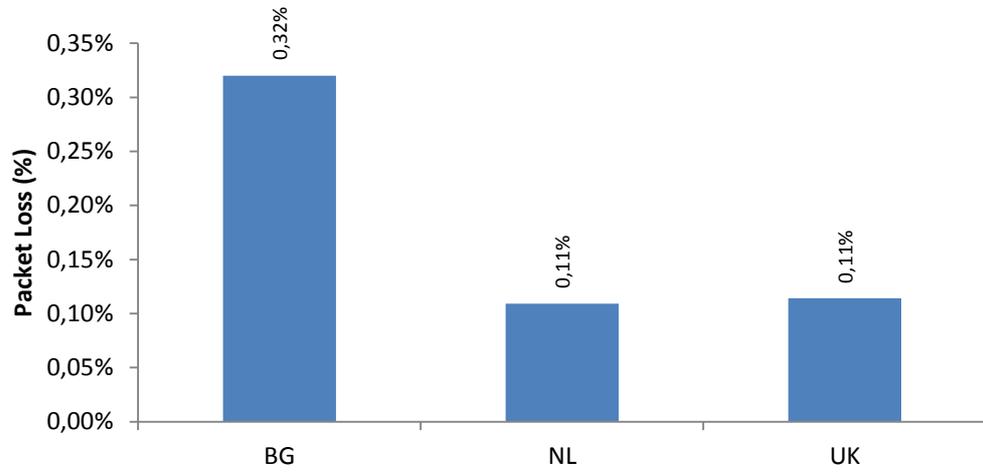


Figure EU.5-17: Weighted Packet loss of FT Tx technology during peak periods, split by country

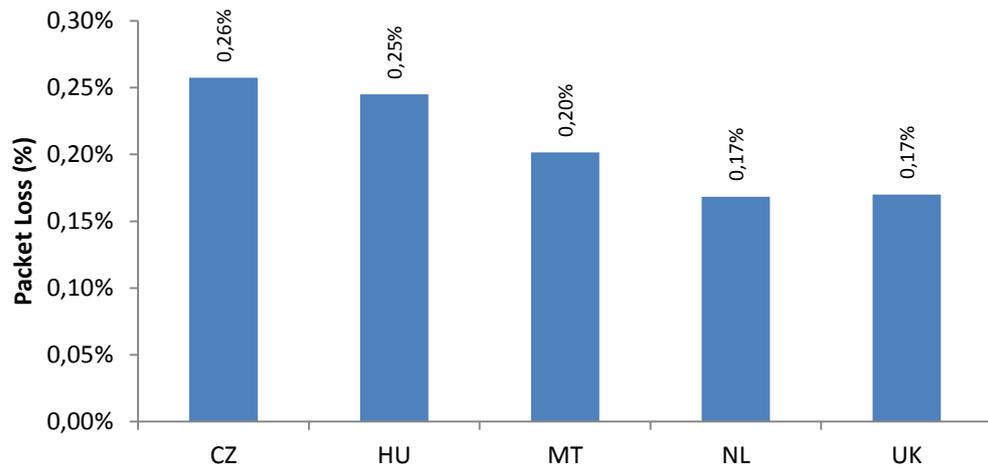


Figure EU.5-18: Weighted Packet loss of cable technology during peak periods, by country

G.2.5 DNS Resolution Time and Failure Rate

Figures EU.5-19 to EU.5-21 show DNS resolution time for all technologies split by country during peak hours. FTTx technology delivers the best DNS resolution times across all countries and technologies apart from the UK. This is due to the Netherlands and Bulgaria deploying FTTH technology. Both countries outperform all countries using cable technology as well. DNS resolution times of xDSL technology across all countries are considerably larger than resolution times of other technologies.

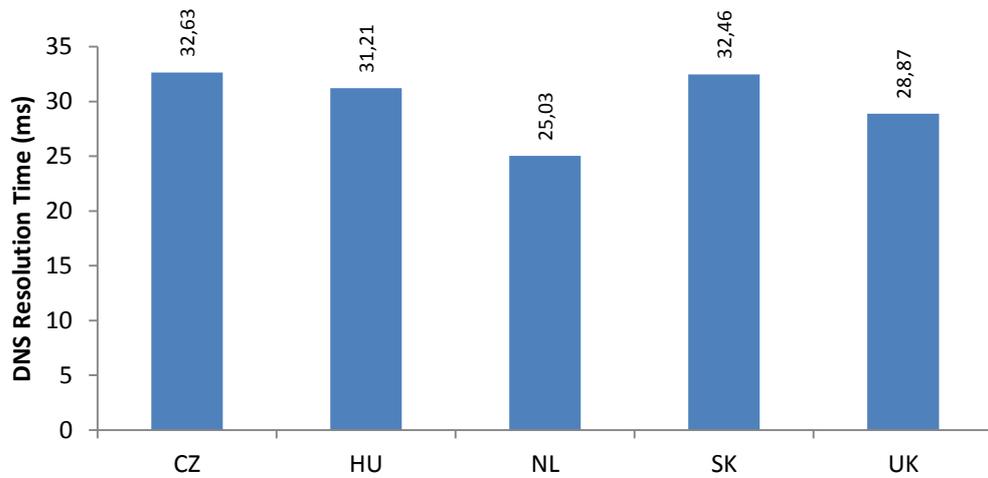


Figure EU.5-19: Weighted DNS Resolution Time of xDSL technology during peak periods, by country

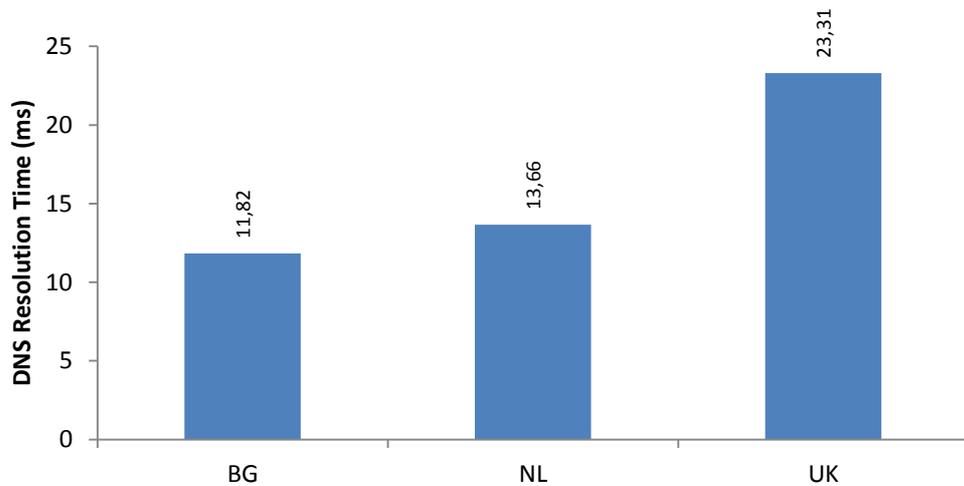


Figure EU.4-20: Weighted DNS Resolution Time of FT Tx technology during peak periods, by country

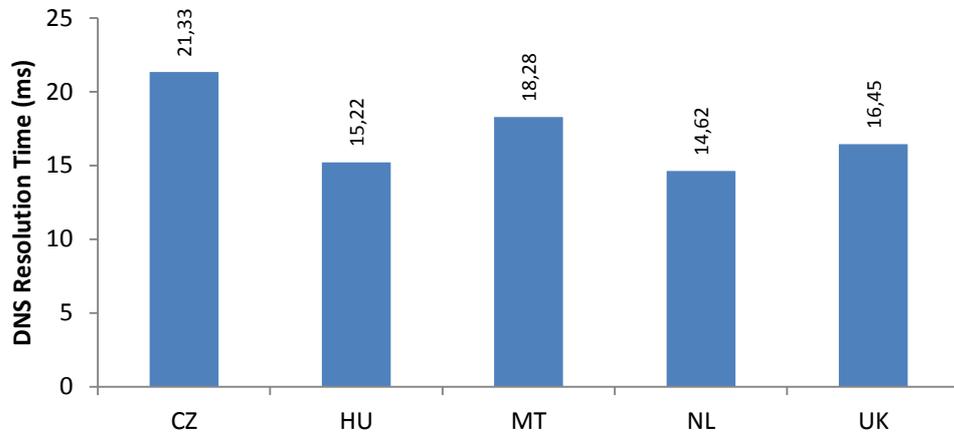


Figure EU.4-21: Weighted DNS Resolution Time of cable technology during peak periods, by country

Figures EU.5-22 to EU.5-24 show DNS failure rate per country for each access technology during the peak period. Most countries using cable and FTTx technologies display very low failure rates below 0.2%, with the Netherlands consistently remaining at or below 0.17% across all technologies. FTTx technology in these countries performs better than xDSL and cable. Contrasting with latency and DNS results, DNS failure rates in Bulgaria are relatively high, particularly compared to the Netherlands and the UK, with the UK's failure rates being very low compared to its resolution time particularly for FTTx and cable.

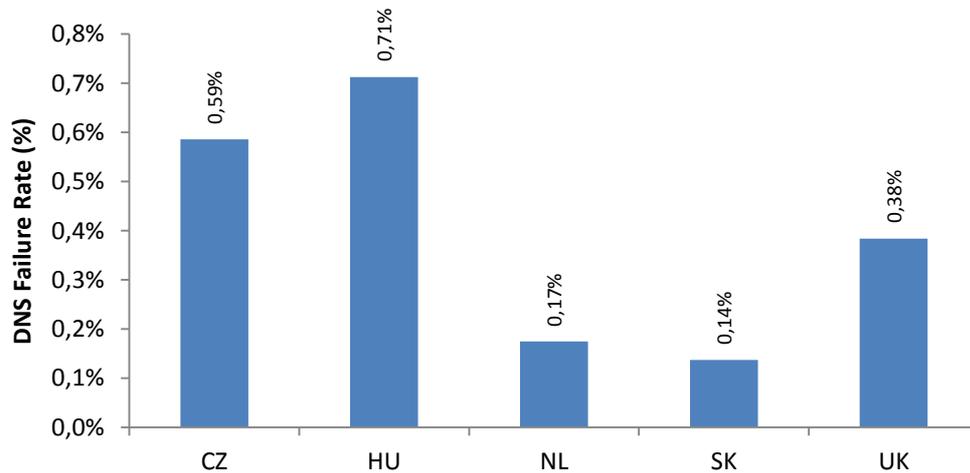


Figure EU.5-22: Weighted DNS Resolution Failure Rate of xDSL technology during peak periods, by country

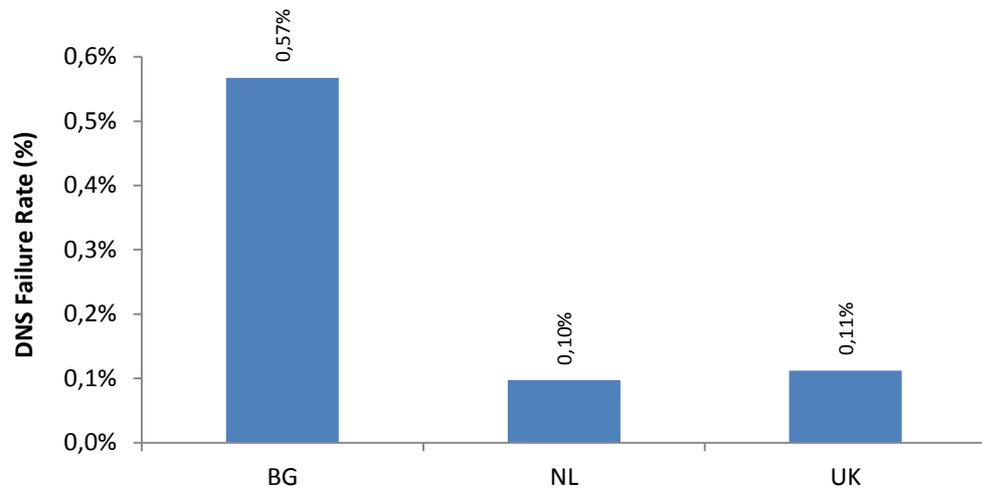


Figure EU.5-23: Weighted DNS Resolution Failure Rate of FTTx technology during peak periods, by country

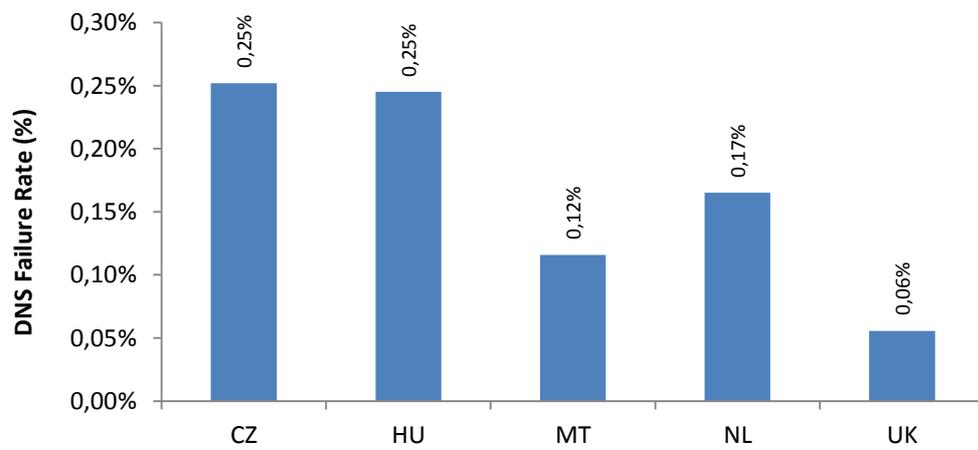


Figure EU.5-24: Weighted DNS Resolution Failure Rate of cable technology during peak periods, by country

G.2.6 Web Browsing Speeds

Figures EU.5-25 to EU.5-27 below display webpage loading times in each country during the peak period, split by access technologies. The test was performed to the public-facing websites of Facebook, Google and YouTube, whose servers are geographically hosted across Europe to optimize consumer performance.

Webpage loading time is between 1 and 3 seconds for most countries using xDSL technology. Loading times are consistently below 1 second for countries making use of Cable and FTTx technologies with the exception Malta for cable technology. This matches the behaviour of latency of this countries as well as reflecting the behaviour of throughput in the downstream and upstream directions. This relationship is also generally observed across all countries and technologies, with Bulgaria proving to be the exception, displaying a high webpage loading time for FTTx despite its very low latency.

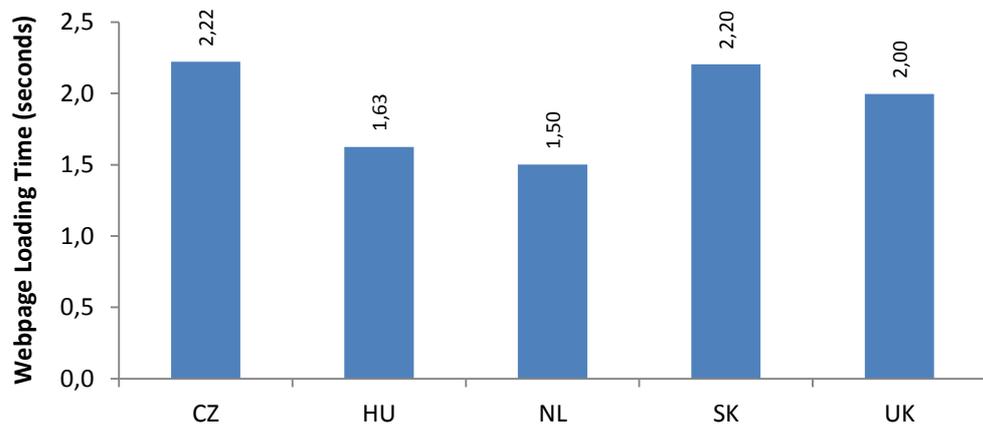


Figure EU.5-25: Weighted Webpage Loading Time of xDSL technology during peak periods, by country

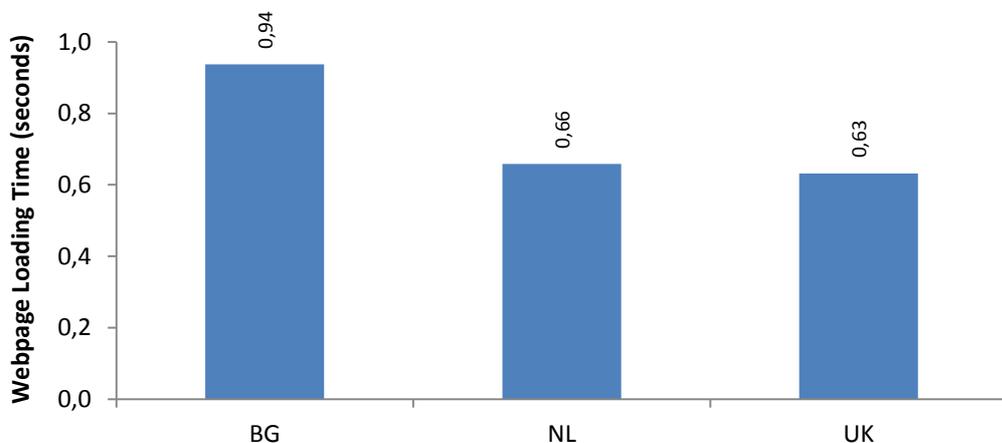


Figure EU.5-26: Weighted Webpage Loading Time of FTTx technology during peak periods, by country

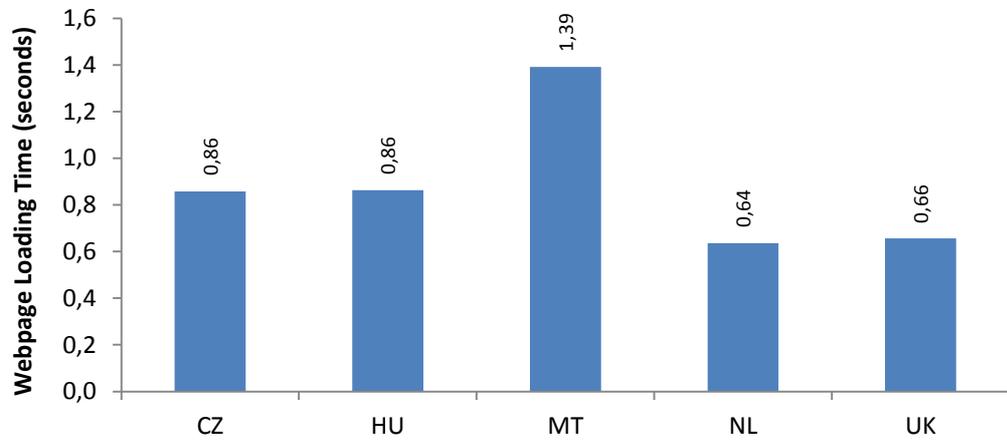


Figure EU.5-27: Weighted Webpage Loading Time of cable technology during peak periods, by country

G.2.7 VoIP Jitter

Downstream jitter is shown in figures EU.5-28 to EU.5-30 during the peak period for each access technology, split by country. Jitter in the downstream direction is very low for most countries across all access technologies. The behaviour of downstream jitter across all technologies generally matches that of webpage loading times, with xDSL technology in Slovakia and the UK being some exceptions. Downstream jitter in Malta and the UK also differ from webpage loading time, with the former performing in line with other countries, contrary to the latter. However, downstream jitter is generally very low, with xDSL technology in Slovakia and FTTx technology in Bulgaria meeting or exceeding 1ms. All other countries display jitter of under 1ms across all access technologies.

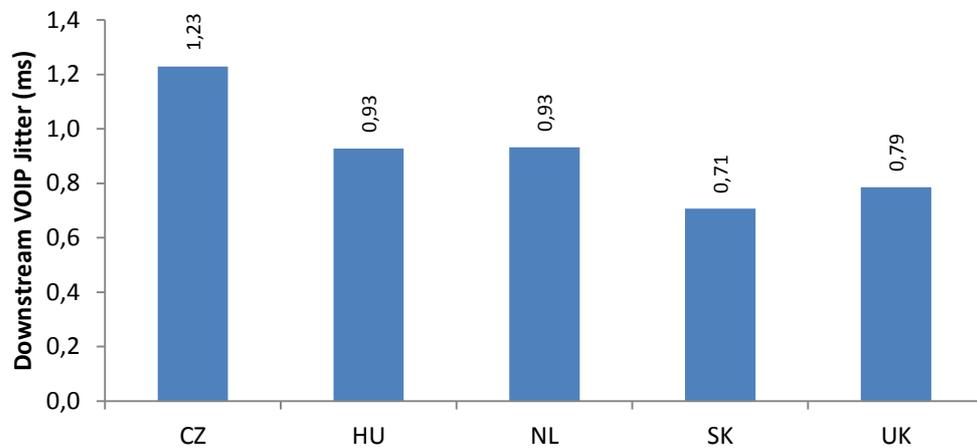


Figure EU.5-28: Weighted Downstream VoIP Jitter of xDSL technology during peak periods, by country

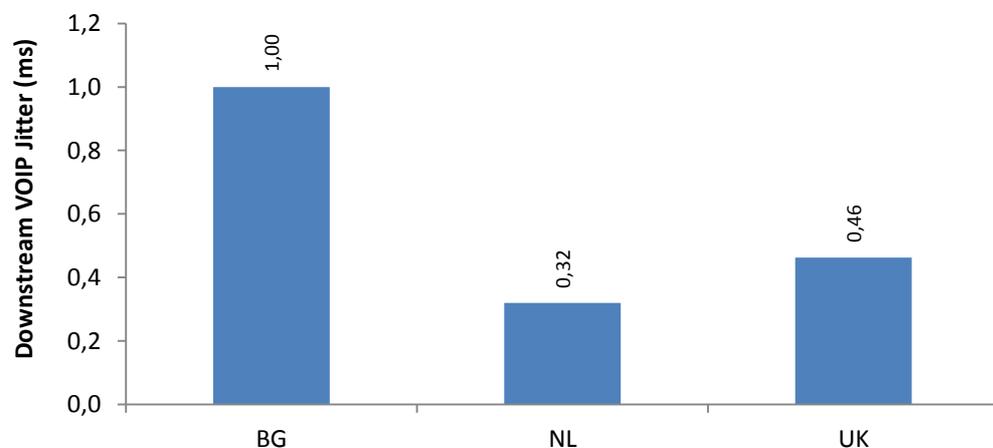


Figure EU.5-29: Weighted Downstream VoIP Jitter of FTTx technology during peak periods, by country

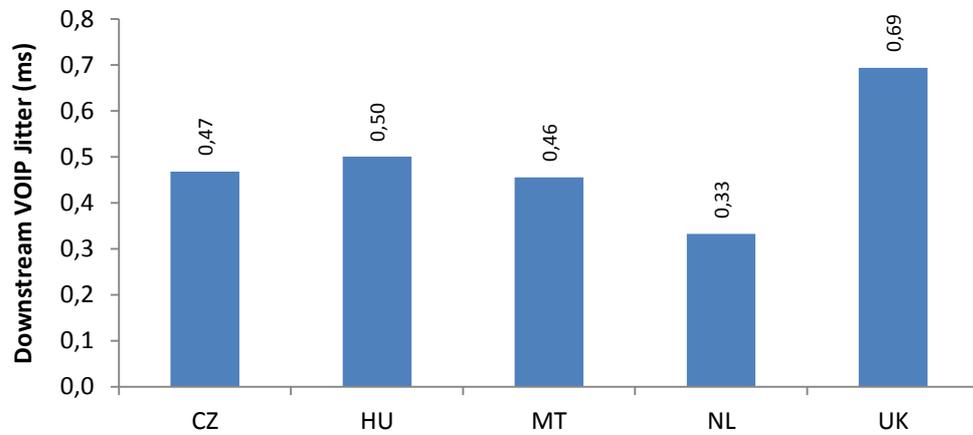


Figure EU.5-30: Weighted Downstream VoIP Jitter of cable technology during peak periods, by country

Figures EU.5-31 to EU.5-33 show results for upstream jitter during the peak period for all access technologies, split by country. Upstream jitter performance also greatly resembles that of webpage loading time across all countries and technologies. As seen in previous sections of this study, upstream jitter of cable technology is higher across all countries compared to xDSL and FTTx technologies. Only Hungary displays a slightly higher upstream jitter for xDSL technology compared to cable. Cable technology in Malta displays the highest upstream jitter across all countries and technologies.

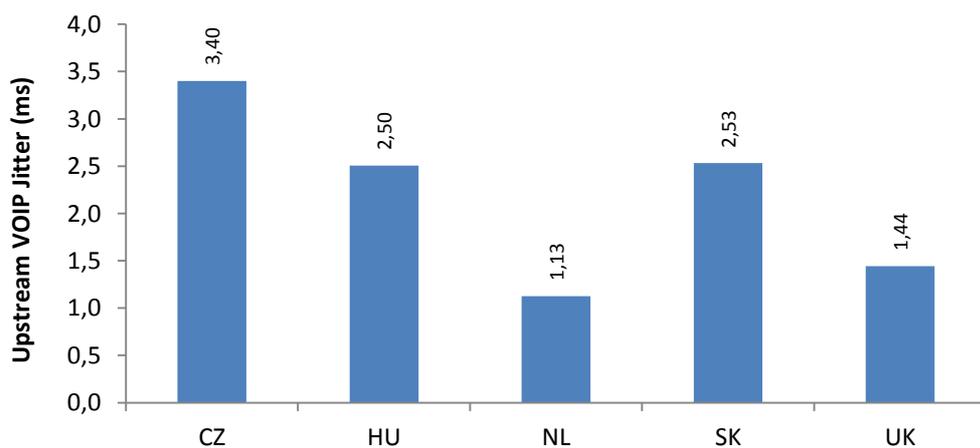


Figure EU.5-31: Weighted Upstream VoIP Jitter of xDSL technology during peak periods, split by country

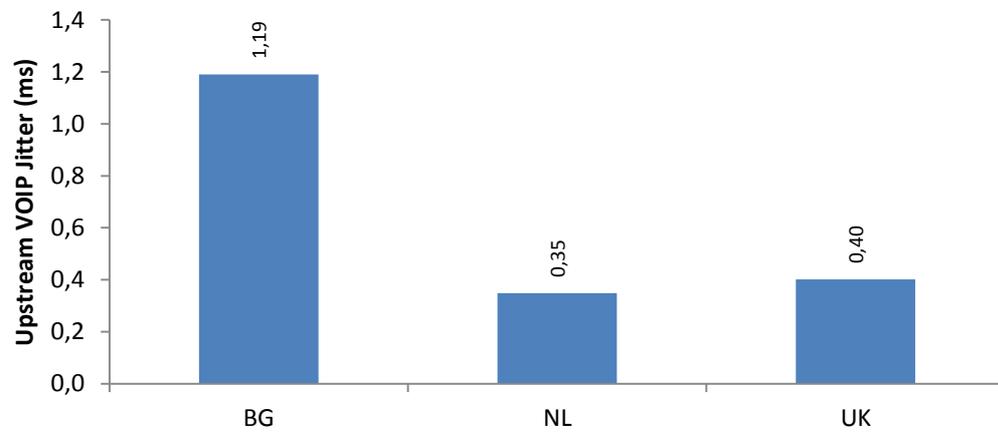


Figure EU.5-32: Weighted Upstream VoIP Jitter of FTTx technology during peak periods, by country

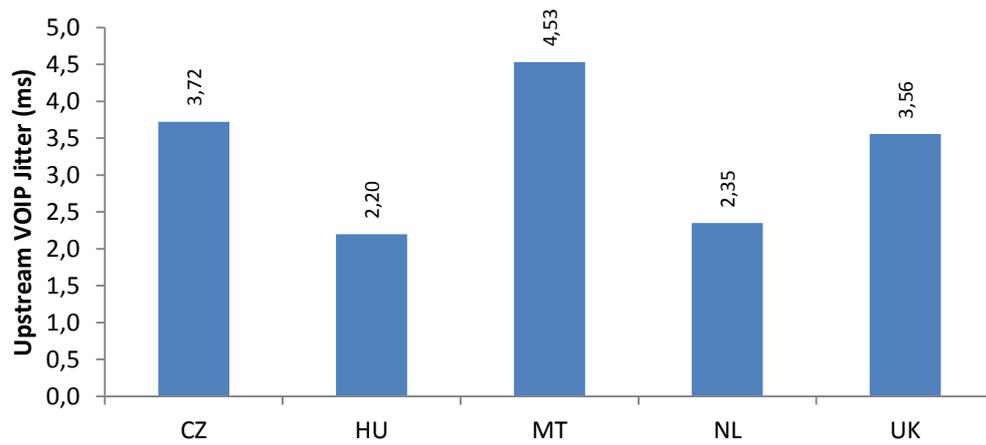


Figure EU.5-33: Weighted Upstream VoIP Jitter of cable technology during peak periods, by country

European Commission

Luxembourg, Publications Office of the European Union

2015 - 188 pages

ISBN 978-92-79-49486-4

DOI 10.2759/770879

