


UK home broadband performance

Annexes

Published 14 September 2023



Contents

Section

A1.	Technical annex	3
A2.	Statistical annex	9
A3.	Sources.....	19
A4.	Glossary of terms	22

A1. Technical annex

Technical methodology

This report is Ofcom's twenty second fixed-line residential broadband speeds report and the nineteenth in which we have published broadband provider package-specific data and comparisons between broadband providers. The technical methodology chosen is the same as that used in Ofcom's previous reports and is based on that created by broadband performance company SamKnows Limited, Ofcom's technical partner in this research project.

SamKnows recruited a panel of UK residential broadband users and supplied monitoring units, known as 'Whiteboxes', to each panellist. SamKnows also managed the collection and aggregation of the performance data and assisted Ofcom in the analysis of the data.

All panellists were sent a hardware monitoring unit which they were instructed to connect to their router. The monitoring unit sits between the panellist's router and the rest of their network, thereby allowing the unit to determine when the network is free to run tests (the device operates in a bridging mode, rather than routing).

The measurement units are connected to panellists' routers using an Ethernet cable in order that the test results accurately reflect the performance of their connections. Where people use Wi-Fi (or other technologies such as powerline) to connect devices to their router, it is possible that the actual speeds received will be lower than those delivered over an Ethernet connection because of the limitations of these technologies. The potential for this difference is greater for higher-speed broadband connections, where the speeds delivered may be higher than the maximum bandwidth that the in-home network technology can support.

This report also used broadband provider-provided data where the SamKnows Router software development kit (SDK) had been embedded directly into the customers' CPE (Customer Premises Equipment). The Router SDK includes all the SamKnows quality of service (QoS) and quality of experience (QoE) tests, as well as cross-traffic detection capabilities.

For the second time, this report features results from a much larger dataset of CPE-embedded tests, which were collected by SamKnows as well as ASSIA Inc. Data from this larger dataset is clearly marked throughout the reports. The test methodology used is the same as that employed in the 'Whitebox' and as described below.

SamKnows developed a customised OpenWRT firmware image which is installed on the units. At the point of delivery to the panellists, this is all that is present on the device; the physical unit contains no additional software, apart from a single script that checks for the availability of the software component at boot-up. This is beneficial both from a security perspective (everything is destroyed when the power is lost) and from a support perspective (any problems with a unit's configuration can be undone simply by power-cycling it). New versions of the software can be delivered remotely without requiring a reboot.

Software within the unit then performed a range of tests to a set schedule, with individual panellists running up to 82,000 (generally around 73,000) separate tests over the course of a day. The software was configured to identify other network activity and not to run tests when such activity was detected. This avoided compromising results by running tests at a time when bandwidth was being used by other internet-connected devices in the household (including those using a wireless connection).

The software uses a combination of standard UNIX tools and custom code developed in C and C++.

All monitoring units maintain accurate time using Network Time Protocol (NTP).

We believe that this technical methodology is robust as it does not rely on monitoring solutions that do not account for the impact on speed of PC set-up, or for having more than one computer using a broadband connection.

Speed tests

The project uses speed tests with multiple concurrent transmission control protocol (TCP) connections to assess the capacity of the user's broadband connection.

Speed tests run for a fixed duration of 10 seconds if the user's broadband connection is not subject to a data cap or has had it lifted for the purposes of this project. Fixed-duration speed tests ensure comparability across broadband connections regardless of their access speed.

On connections slower than 30 Mbit/s, units download 6 MB using separate TCP sessions (in parallel) or transfer for 10 seconds (whichever is reached first). Connections faster than 30 Mbit/s but slower than 50 Mbit/s will transfer an increased amount during the downstream throughput test. This amount is up to 12 MB or 10 seconds (whichever is reached first). Connections of 50 Mbit/s or faster are all without data caps and therefore employ the full 10-second speed test.

The nature of the protocols used on the internet means that during a file download the speed at which data is sent is gradually increased until a stable speed is achieved. To measure this stable speed, our tests exclude the period of the speed ramp-up. The exact way in which the speed ramp-up occurs on different networks may lead to slight variations in the accuracy with which the stable speed can be measured.

An initial lead-in period is used to ensure that TCP window sizes are increased before measurements are made. Multi-thread tests were run 12 times per day, once every three hours in off-peak periods and once every hour at peak times. We found that, typically, the download speeds achieved using the multi-thread tests in the early hours of the day determine the maximum speed the line can support.

Additionally, it is understood that some broadband providers operate transparent HTTP proxy servers on their networks. To overcome this, the web servers are configured to respond with the following headers, which should disable caching in standards-compliant proxy servers:

```
Cache-Control: "private, pre-check=0, post-check=0, max-age=0"
```

```
Expires: 0
```

```
Pragma: no-cache
```

As with the download tests, connections with upload speeds slower than 10 Mbit/s upload 3MB or transfer data for 10 seconds (whichever is reached first). Connections with upload speeds higher than 10 Mbit/s and lower than 20 Mbit/s upload 6MB or transfer data for 10 seconds (whichever is reached first). Connections with upload speeds higher than 20 Mbit/s employ the full 10-second speed test. Four speed-test servers are deployed in data centres in London to handle the traffic. 40 Gbit/s of capacity is shared between these servers. Each server is monitored for excessive network load and for CPU, disk and memory load.

The test results gathered by each server are compared against one another daily, to ensure that there is no significant variation in the speed attainable per server. Units cycle through the speed-test servers in a round-robin fashion when testing.

Testing web page loading times

The test downloaded the HTML and media assets of a simple web page hosted on a SamKnows-managed server. This makes use of up to eight concurrent TCP connections to fetch the assets. Both tests make use of libcurl.

The time in milliseconds to receive the complete response from the web server is recorded, as well as any failed attempts. A failed attempt is deemed to be one where the web server cannot be reached, or where a HTTP status code of something other than 213 is encountered.

Tests were run every hour.

Testing latency, packet loss and jitter

A bespoke application was used to test latency and packet loss. The application was designed to run continuously to get a statistically robust set of data. The test used UDP rather than ICMP and sent approximately 2000 packets every hour.

Some broadband connections suffer from excessive latency and packet loss when the connection is being heavily used to download or upload data. In normal operation, the SamKnows UDP latency and packet loss test is paused during other tests, to ensure that the latency and loss statistics are not adversely impacted by these tests. When configured to report the latency-under-load metric, the UDP latency and loss test sends and receives packets during our download and upload speed tests, reporting latency and loss statistics (as 'latency under load') after each test as a separate metric, depending on the direction of the test.

The test also records instances of contiguous packet loss events. These are termed 'disconnections'. The duration of the disconnection event will vary by its cause – a minor routing issue may cause only a few seconds' disconnection, whereas a modem losing synchronisation with the telephone exchange may result in a 30-second disconnection.

Another bespoke application was used to test jitter. This application sends 500 packets upstream and 500 packets downstream via UDP, every hour, and records the number of packets it sent and received (thereby providing a loss rate), and the jitter observed for packets it received from the server. The server does the same, but with the reverse traffic flow, thus providing bi-directional loss and jitter.

Testing recursive DNS resolver responsiveness and failures

Testing a broadband provider's recursive DNS resolution can be accomplished using many tools, such as nslookup, dnsip and dig. A custom DNS measurement client is used so that we can support regular UDP-based DNS queries, as well as emerging standards such as DNS-over-HTTPS and DNS-over-TLS.

For regular UDP-based DNS queries, as used in this report, the Whitebox will honour the DNS servers offered by the home router over DHCP. This router may provide the DNS servers of the broadband provider directly, or it may provide its own IP address and then proxy the DNS requests upstream. In all cases, we record the DNS server IP address that the router provided over DHCP.

The tests record the number of milliseconds for a successful result to be returned. A successful result is deemed to be one when an IP address is returned (the validity of the IP address is not checked). A failure is recorded whenever the DNS server could not be reached, or an IP address was not returned. The hostnames of four popular websites were queried every hour.

Testing Netflix video streaming performance

The Netflix test is an application-specific test, supporting the streaming of binary data from Netflix's servers using the same content delivery network selection logic that their real client uses. The test has been developed in direct cooperation with Netflix.

The test begins by calling a Netflix-hosted web-based API. This API examines the client's source IP address and uses the existing proprietary internal Netflix logic to determine from which Netflix server this user's IP address would normally be served content. This logic will consider the broadband provider and geographic location of the requesting IP address. Where the broadband provider participates in Netflix's Open Connect programme, it is likely that one of these servers will be used. The API will return to the client a HTTP 302 redirect to a 25MB binary file hosted on the applicable content server.

The test then establishes a HTTP connection to the returned server and attempts to fetch the 25MB binary file. This runs for a fixed 20 seconds of real time. HTTP pipelining is used to request multiple copies of the 25MB binary, ensuring that if the payload is exhausted before the 10 seconds are complete, we can continue receiving more data. The client downloads data at full rate throughout; no client-side throttling takes place.

It is important to note that this 25MB binary content does not contain video or audio; it is just random binary data. However, with knowledge of the bitrates that Netflix streams content at, we can treat the binary as if it were video/audio content operating at a fixed rate. This allows us to determine the amount of data consumed for each frame of video (at a set bitrate) and the duration that it represents. Using this, we can then infer when a stall occurred (by examining when our simulated video stream has fallen behind real time). The test currently simulates videos at bitrates of 235 Kbit/s, 375 Kbit/s, 560 Kbit/s, 750 Kbit/s, 1,050 Kbit/s, 1,750 Kbit/s, 2,350 Kbit/s, 3,000 Kbit/s, 4,500 Kbit/s, 6,000 Kbit/s and 15,600 Kbit/s.

The test captures the 'bitrate reliably streamed' (the highest quality video that can be streamed without rebuffering events), the download speed from the Netflix Open Connect Appliance and the video start-up delay.

Connections with usage caps

Some of the test units were deployed on broadband connections with relatively low usage caps. To avoid using a significant proportion of the available download limit each month, the test schedule for the test units on these connections was reduced.

Research methodology

The performance data in this report are taken from a base of 11,901 panellists. This includes data from 2,238 panellists on SamKnows' independent global platform, SamKnows One, who had a broadband monitoring unit connected to their routers in March 2023, 407 panellists who had the SamKnows SDK embedded into their CPE and 9,256 panellists from an extended embedded data collection from broadband monitoring units connected as part of several ISP's private panels. This extended panel was used for broadband provider package performance comparisons. Ofcom's definitions of geographic broadband markets use the definitions for the wholesale broadband access (WBA) market review. These were an important consideration in recruiting our panel and applying statistical analysis, because they enabled us to ensure that our panel was representative of the UK residential broadband market overall, and facilitated like-for-like comparison between broadband provider packages:

- Each panellist was assigned to one of the geographic markets, and we weighted the analysis accordingly to ensure that our overall findings were representative of UK residential broadband performance overall (for example, as Market A represents around 1% of UK premises, we ensured that performance data from panellists in Market A contributed 1% towards the overall computation of UK residential broadband performance).
- For comparisons of broadband provider package performance, we used only panellists who live within geographic Market B. This means that all panellists used for the broadband provider package comparisons live in areas served by a local telephone exchange in which at least one operator other than BT is present, i.e. there is at least one local loop unbundling (LLU) operator. This avoids any potential distortions of the data by broadband providers using BT wholesale services (BT Retail, EE and Plusnet), caused by the inclusion of panellists who live in (typically less densely populated) Market A areas, and to whom LLU services are not available.

We have used statistical techniques to adjust our results to ensure that they are representative of the overall UK broadband population. This includes weighting the results from our panel by rural/urban, distance from exchange, geographic market definition and broadband provider. For the provider-specific comparisons we have also ‘normalised’ the data for ADSL operators by distance from exchange (using the straight-line distance from the panellist’s location to the exchange), which we believe is necessary to provide like-for-like comparisons of broadband providers which have different customer profiles.

David Saville of Saville Rossiter-Base has assessed the research methodology and panel and helped ensure its suitability for purpose. Checks were also applied to ensure that straight-line distance was an appropriate metric for carrying out normalisation, including comparing this distance with the line attenuation. Details of the statistical methodology used are provided in Annex 2. The methods of analysis for the provider-specific comparison are based on those used in the July 2009 report which had expert review by econometrician Professor Andrew Chesher of University College London.

Router-embedded testing

Broadband performance test results from the extended embedded data collection included in this report were collected by the participating ISPs (BT, Virgin Media, TalkTalk and Sky) hourly between 3-5am as well as 8-10pm for the month of March 2023. BT, Virgin Media, TalkTalk and Sky collected download speed, upload speed and latency measurements. We requested that the ISPs adhere to the following testing principles:

1. Each panellist must have a unit capable of running the appropriate test software that can measure download and upload speeds as well as latency received at the customer premises equipment (CPE).
2. The tests should not be run when user traffic is detected by the unit, as this could result in a negative impact on the performance experienced by the user, and may also compromise the test results.
3. To avoid detriment to panellists, the data used for the tests must not be included in the panellists’ data allowance (if any).
4. To assess the capacity of the user’s broadband connection, up to three concurrent transmission control protocol (TCP) connections must be used.
5. The download and upload tests will consist of downloading and uploading files over a duration of at least 5 seconds starting from the end of the ramp-up period. (When a

measurement is being conducted, the results must not be taken before the ramp-up period has been completed, and a stable speed has been reached. The ramp-up period is considered complete when three consecutive transfers of small blocks of data of the payload of the target file are downloaded at the same speed or within 15% of one another. At this point the results can be taken).

6. The application layer protocol must be http 1.1.
7. Actual speeds must be sampled from the CPE (or the ISP's test equipment) in the end user's premises through to the internet interconnect, which must be a representative independent internet exchange point.
8. The server should not be located on the ISP's network, and traffic should ideally be routed to a London test server. We expect ISPs to procure their own server for the testing.

A2. Statistical annex

Key statistical concepts used in this report

This report presents the findings from research which has involved the collection and interpretation of 29.7 million data points. It has been a complex process, both technically and statistically.

The glossary in Annex 4 provides definitions of the technical terms we use throughout the report. However, knowledge of the following is important to understand how we have analysed the performance data collected.

- We present data in the report only in cases where there are sufficient data points to deliver a statistically sound result. This means that we report performance only when statistical analysis indicates that our findings are accurate enough to be useful. Accuracy is determined by the number of measurement tests undertaken, the size of the sample (number of panellists) and the variation (spread or range of results) between panellists.
- In order to acknowledge the limited accuracy of the estimates, and to ensure that we highlight only those differences that are statistically significant, for many charts we do not show a value but instead show a range around the mean value which indicates the statistical confidence we have in our results. The range we use is called a 95% confidence interval, which is a statistically derived range calculated from the standard error (which is itself calculated from the sample size and the variation within the sample). A 95% confidence interval means that if we repeated the research with a different sample, assembled in the same way, there would be a 95% probability that the mean value would be in the range shown. Where we have large samples and/or little variation within the sample, the confidence interval is much narrower than where we have smaller samples and/or large variation within the sample. Differences are reported as significant if they are significantly different as judged by a two-tailed 5% test of statistical significance. In the tables where we present differences that are statistically significant, we present differences that are significant to a 95% level of confidence and highlight those that are significant to a 99% level of confidence by using an asterisk.

Similarly, when we have used the median as a measure of central tendency, we have calculated 95% confidence intervals for the median value. Due to the skewed distribution of several measures, we have also used the median as a measure of central tendency throughout the report and calculated 95% confidence intervals around this value. For a given metric, if we have n units and place every unit's value on this metric in ascending order, every unit will have a position from 1 to n . The unit with the lowest value of this metric will have a position of 1, and the unit with the highest value will have a position of n . To find the median, we find the unit with a position of $(n+1)/2$ in the ordered dataset and then find the value that this unit has on the metric. This leaves 0.5 of the dataset with a value below the median and 0.5 with a value on this metric higher than the median. Then, to calculate the 95% confidence interval around the median, we used a calculation that uses a normal approximation to the binomial distribution and returns the position of the value that is the upper bound of the median and the position of the value that is the lower bound of the median. To get the upper and lower bounds, refer to the ordered dataset and use the values at the relevant upper and lower bound positions. While the upper and lower bound positions are always an equal distance from the median position and closer to the median position in larger datasets, as the corresponding values relate to actual metric values obtained by a unit in the dataset, sometimes the upper and lower bound values are different distances from the median values. Also, the analysis in this report

has been conducted on weighted data. With weighted data, the calculation for the median and upper and lower bounds is the same; however, rather than there being n units, with each counting for one position, the value of n relates to the sum of all the units multiplied by their weight, and the positions relate to the cumulative sum of weights at each point in the data set. The table below shows the distinction between how we find the median and its upper and lower bounds for weighted and unweighted data.

Figure 2.1: Finding the median for weighted and unweighted data

Absolute Position	Metric Value	Weight	Weighted Position	Unweighted Data	Weighted Data
1	60	2	2		Lower Bound
2	112	1	3	Lower bound	Median
3	130	0.5	3.5	Median value	
4	132	0.5	4	Upper bound	Upper bound
5	147	1	5		

Source: Ofcom

- In order to ensure that the national data we present are representative of UK residential broadband users as a whole, we have weighted the data by broadband package, nation and rural/urban split, market classification, distance from the exchange (for ADSL packages) and maximum attainable speed (for FTTC packages).
- We have similarly weighted the data where we are comparing the performance of individual broadband packages, in order to ensure that the analysis provides a fair comparison of actual performance rather than reflecting random differences in package customer profiles in the sample.
- To ensure that comparisons between the performance of different technologies in urban and rural areas are representative of the urban and rural markets, we applied a separate weight which gave a representative sample by broadband package, nation, market classification technology, speed tier and urbanity when making these comparisons. This weight is a change in the methodology used in previous years, where the weight only considered technology, speed tier and urbanity.
- A difficulty in comparing ADSL and FTTC broadband providers is that with this technology, speed varies by the length and quality of the specific consumer's telephone line. Therefore, providers which have a higher proportion of customers in rural areas, where line lengths are typically longer, may be expected to deliver lower speeds on average than those which focus on towns and cities, simply because they have a different customer profile. For FTTC customers, the critical part of the line is that between the customer's house and the cabinet – this section of the line is copper and subject to line degradation.
- To resolve this issue, we have taken the following steps:

- > For all broadband provider package comparisons, we have included only consumers who live in an area where the exchange has been ‘unbundled’ by at least one LLU operator. This means that broadband providers using wholesale services (such as BT Wholesale’s IP stream or Wholesale Broadband Connect products) can be compared on a like-for-like basis with LLU operators.
- > We have excluded all ADSL customers where the straight-line distance from their home to the local telephone exchange is more than 5km, to limit the impact of outliers when weighting, and we have normalised data to straight-line distance distributions.
- > Straight-line distance weighting was applied only to ADSL operators in this report and not to cable or FTTC services, where performance is less influenced by distance from the exchange.
- > For FTTC customers, we do not have adequate information on distance between cabinet and customer premises. We therefore approximate this distance by normalising data using the maximum attainable speed. The maximum attainable speed is the best speed which a line can carry and is therefore a suitable proxy for measuring the quality of the line.
- > No weights were applied to Virgin Media cable packages or FTTP packages, as these circuits are not affected by distance from the exchange or supplier cabinets.

Sample methodology

A panel of UK residential broadband users was drawn from a pool of volunteers following a recruitment campaign by SamKnows. The objective was to obtain a representative panel, in order to monitor the performance of residential fixed-line broadband in the UK over a two-year period of research. In addition to obtaining a panel sufficient for monitoring changes in overall performance, the panel was recruited to enable specific analysis of the performance of the most common broadband packages in the UK, in particular higher-speed packages.

The main purposes of this recruitment have been to:

- replace panellists who leave the panel due to natural attrition, such as moving to a new house or losing interest in participating in the research;
- ensure adequate samples for all broadband providers and replacing panellists who, although remaining on the panel, have decided to switch their operator and/or package. As the Ofcom panel tends to be comprised of people with an interest in telecoms, there is a strong focus on ensuring coverage of lower speeds packages; and
- enable the broadband provider-level reporting of new packages such as high-speed fibre and cable as soon as sufficient panellists can be recruited.

Due to shortfalls in some areas, SamKnows provided additional data from its independent global platform, SamKnows One. At the moment, the total active panel is 2,557 ‘active’, meaning those who contribute results to the national analysis. 2,150 of these respondents belong to the SamKnows’ panel and 407 to SamKnows’ enabled embedded panel. Additionally, a further 9,256 units from ISP’s private embedded panels were used in the broadband package comparisons. However, data from these units were not included in the national or rural analysis.

Before they received the monitoring units, volunteers were screened and preliminary speed measurements and checks on their IP addresses were undertaken, to reduce the impact of respondents' misconceptions during sampling regarding which package they were using.

Definition of valid panellists and test volumes

All measurement data were collated and stored for analysis purposes as a monthly trimmed average of the measurements obtained for each respondent for the relevant time interval (e.g. 24 hours, 8pm to 10pm weekdays, 9am to 5pm Monday to Friday). Only panellists who provided a minimum of five valid measurements across all the download speeds, upload, latency, DNS and webpage tests for the narrow 8pm to 10pm peak measurement period, and a good spread of tests throughout the day, were included in the monthly analysis. A trimmed mean was used because, for a small proportion of respondents, the occasional test result was far more than what was achievable on the line. The top and bottom 1% of results per respondents did not count towards the average.

The average number of measurements per respondent for the 24-hour multi-thread download speed tests within the national panel between 1 March 2023 and 31 March 2023 was 386 overall, 319 for those on the SamKnows Panel, from a theoretical maximum of 360 per respondent, and an average of 738 for those on the embedded panel, from a theoretical maximum of 720 per respondent (i.e. if all panellists had their monitoring unit connected on 1 March and all scheduled tests were run. Tests were not run when the monitoring unit detected concurrent use of the bandwidth).

Average download speeds are generally very accurately measured, so the main factors limiting the accuracy of the analysis reported here are the number of panellists and the average number of measurements.

Quotas and weightings

Quotas were set before the exact package market shares for operators became available, but results were weighted to be representative at national level. To recruit broadband provider packages to match the specific quota criteria above, and to achieve 100-150 panellists per package, only those broadband packages with more than 250,000 subscribers in total were targeted, although we do include broadband packages with fewer than 250,000 subscribers where we can recruit sufficient panellists, and where we believe a package is important enough to the future development of the market to warrant inclusion in the report.

The results and analysis of the 11,901 SamKnows and ISP's embedded panellists' measurement results were divided into three separate datasets, each weighted to targets.

National panel (over 2 Mbit/s packages): 2,557 panellists. All had at least five valid test measurements across all download tests, with a validated IP address, single measurement speed check, and distance and geographic market classification data.

Broadband provider package panel: 9,256 panellists. Respondents from both SamKnows' panel and externally recruited embedded data was used for this panel. Respondents for this panel consist of panellists from geographic markets 2 and 3 only. There was a target of 100 valid panellists for each package, but the criterion for inclusion in the reporting was an effective sample minimum of c.50 valid panellists (those with a base of fewer than 75 should be treated with caution). Additional validation for the broadband package panel included a review of measured speed against straight-line distance from the exchange to the panellist's premises, and a review of outliers. Any package reassignment identified was made to both the broadband provider package panel and the national panel datasets.

Rurality panel: 2,557 panellists. All had at least five valid test measurements across all download tests, with a validated IP address, single measurement speed check, and distance and geographic market classification data. The sample composition of this panel is the same as the national panel.

Sample weighting

National panel:

- Weighting by broadband provider market and package shares by LLU/ non-LLU connections supplied by broadband providers as at March 2023, urban/rural, geographic market classification, ADSL distance to exchange (fitted to UK representative exchange line distribution provided by Openreach Limited) and maximum attainable normalisation for FTTC lines.

Broadband provider package panel:

- Weighting to distance from exchange (those panellists with an unrecorded or straight-line distance to the exchange of more than 5km were excluded).
- ADSL2+ packages were normalised by distance from the exchange, to the aggregated distribution of straight-line distance between premises and exchanges of all panellists on those headline packages.
- FTTC packages were normalised to the appropriate maximum attainable speed curve that matched the headline package speed (36 Mbit/s, 50 Mbit/s or 63-67 Mbit/s).
- Sky's 61 Mbit/s package was weighted to the distribution of maximum attainable speed lines serving this package.
- Cable or FTTP packages are not weighted, as speed of service is not directly related to distance from the exchange.

Rurality panel:

- Weighting by broadband provider market and package shares by LLU/ non-LLU connections supplied by broadband providers as at March 2023, geographic market classification, ADSL distance to exchange (fitted to UK representative exchange line distribution provided by Openreach Limited) and maximum attainable normalisation for FTTC lines, and by package speed tier take-up within technology and within rurality.
- The weighting structure of package speed tier take up by was formulated using a combination of May 2021 data provided to Ofcom by internet service providers, and market shares across technologies. This is a different methodology than has been used previously, and therefore unless explicitly stated, urban rural statistics are not directly comparable to those of previous years.

Assigning panellists to broadband provider and broadband package

The following process was applied, to select panellists and assign them to the correct broadband provider package:

- Volunteer panellists were required to provide their broadband provider, package name, headline speed and download limit from drop-down menus and/or text boxes provided in an online form. This was used as initial categorisation of potential candidates against the target quotas.

- > The stated package name and headline speed (where they allowed identification of the correct broadband provider package) were used to assign panellists to a package.
- Volunteers who matched the sample criteria were screened by broadband provider package, and an average speed reading estimate was obtained to screen actual versus stated package. Those who were successfully screened were sent monitoring units.
 - > The stated broadband provider allocation was validated against IP address. When an IP address and stated broadband provider were inconsistent or missing, the volunteer was rejected. When an average speed measurement was outside the feasible range, the volunteer was flagged, and a monitoring unit box dispatched if there was still sample required for the assessed package.
- Once the volunteer correctly connected the monitoring unit and test measurements were received, straight-line distance from home to exchange and geographic market classification were added to the measurement data.
- A further stage of ensuring that respondents were assigned to the correct broadband provider package took place before the analysis stage. The following steps were undertaken:
 - > The initial assumption was that the package assignment, recorded in the panel data file, was correct. However, the broadband providers were asked to verify that respondents were on the correct package.
 - > However, those participants whose stated and measured package assignments or broadband provider were not consistent, and could not be definitively reconciled, were excluded from the comparison data. Only those panellists with an ADSL connection, who were connected to an ADSL2+ enabled exchange, were considered for an ADSL2+ package allocation. The above modification (upload speed assignment) was necessary to identify those customers using ADSLMax on an ADSL2+ exchange.

Weighting to distance from exchange

As performance of ADSL broadband is significantly affected by the length of the line between a consumer's premises and the local exchange, any comparison between broadband provider or technology could be affected by the distribution of distance among the sample.

It was therefore necessary to weight the data by distance from the exchange in order to provide like-for-like comparison between the previously published data, to ensure that any differences identified were due to differing performance and not due to a differing distribution of line lengths. Openreach Limited provided three curves which indicate the national distance profile of ADSL1, ADSL2+ and all ADSL lines for all lines in the UK. Each relevant ADSL2+ broadband provider package in the broadband provider package panel is adjusted to match this national profile. ADSL packages in the national panel are adjusted to match the profile of all ADSL lines as provided by Openreach Limited.

Distance from premises to local exchange was captured as the straight-line ('as the crow flies') distance, measured from the full postcodes of premises to the local exchange.

Weighting fibre packages

Although fibre technologies show little speed degradation between the local exchange and the final point where fibre is present, most respondents with fibre have FTTC only. This means that the length of the co-axial cable between the cabinet and the consumer premises can have a significant impact

on speed. As the FTTC network is being rolled out into more rural areas, the distribution of distance from the cabinet becomes important, as rural lines tend to be longer than urban.

In a similar manner as weighting for distance from the exchange, for ADSL, Ofcom has decided to normalise for distance from cabinet for FTTC products, to ensure a like-for-like comparison. An identical model to ADSL, based on straight-line distances from the cabinet, is not possible, as the relevant cabinet for many premises will be in the same postcode. Therefore, a proxy for distance from cabinet was used – this is maximum attainable speed. This is a network metric which assesses the line and determines the maximum speed it can carry. Openreach Limited provided the maximum attainable speed for each panellist, and the profile of fibre lines in the UK. Each broadband provider’s respondent profile is adjusted to match the national profile and weighted accordingly to ensure like-for-like comparisons.

Ofcom uses a single curve for each speed, which does not discriminate between respondents with self- and engineer-installed lines.

Weighting efficiency

Overall, against the entire weighting framework, the national panel achieved a weighting efficiency of 39%. Since there are five factors making up the national weights, the weighting efficiency is often lower than the broadband provider weighting. It is also affected by the composition of our panel and can therefore vary year on year. The under-0.5s are primarily driven by the over-representation (against current market shares) both of some FTTC packages, and panellists in rural areas, and Scotland and Wales. The over-2s are driven by market shortfall for some packages and for under-sampling in market A.

Figure 2.2: National panel range of weights

Weights	Count	Percentage
Less than 0.5	1046	44.82%
0.5 to 1	543	23.26%
1 to 1.5	198	8.48%
1.5 to 2	164	7.03%
More than 2	383	16.41%

Source: Ofcom

Overall, against the entire weight frame, the broadband provider package panel achieved a weighting efficiency of 94%. This is because Virgin Media cable and FTTP (included those from the embedded panel) packages are not weighted, as distance from exchange does not impair download speeds.

Figure 2.3: Broadband provider package range of weights

Weights	Count	Percentage
Less than 0.5	50	2.41%
0.5 to 1	1705	82.21%
1 to 1.5	270	13.02%
1.5 to 2	25	1.21%
More than 2	24	1.16%

Source: Ofcom

Figure 2.4: Weighting efficiency, by broadband provider package

ISP	Package Download Speed	ISP Weighting Efficiency	ISP	Package Download Speed	ISP Weighting Efficiency
Virgin	52	100%	EE FTTP	67	100%
Virgin	125	100%	EE FTTP	300	100%
Virgin	250	100%	KCOM FTTP	50	100%
Virgin	350	100%	KCOM FTTP	75	100%
Virgin	500	100%	KCOM FTTP	100	100%
Virgin	1000	100%	KCOM FTTP	175	100%
BT FTTC	38	59%	KCOM FTTP	250	100%
BT FTTC	52	72%	KCOM FTTP	400	100%
BT FTTC	76	95%	KCOM FTTP	750	100%
EE FTTC	38	94%	Plusnet FTTP	330	100%
EE FTTC	76	95%	Sky FTTP	59	100%
Plusnet FTTC	38	84%	Sky FTTP	145	100%
Plusnet FTTC	76	100%	Sky FTTP	500	100%
Sky FTTC	76	98%	TalkTalk FTTP	110	100%
TalkTalk FTTC	38	86%	TalkTalk FTTP	150	100%
TalkTalk FTTC	76	94%	TalkTalk FTTP	220	100%
BT FTTP	50	100%	TalkTalk FTTP	500	100%
BT FTTP	75	100%	TalkTalk FTTP	1000	100%
BT FTTP	150	100%			
BT FTTP	300	100%			
BT FTTP	330	100%			

Source: Ofcom

Overall, the rurality panel achieved a weighting efficiency of 49%.

Figure 2.5: Rurality panel range of weights

Weights	Count	Percentage
Less than 0.5	883	35.23%
0.5 to 1	774	30.89%
1 to 1.5	257	10.26%
1.5 to 2	257	12.49%
More than 2	279	11.13%

Source: Ofcom

Comparison of urban and rural speeds

Using Bluewave Geographic’s Locale dataset, it is possible to segment all UK postcodes into one of seven urban-rural groupings. This dataset, widely used in market research design and sampling, allocates postcodes to a category based on their population density and how close the settlement is to a larger one. The seven groupings range from A (large cities such as London and Birmingham) to G (isolated rural areas such as the Western Isles and Dartmoor).

To simplify the analysis, the groupings have been banded together into two broad groups: urban and rural (population less than 2,500 and in open countryside). This grouping enables us to compare rural and urban areas over time.

Additional embedded data

Broadband performance test results from the extended embedded data collection included in this report were collected by the participating ISPs (BT, Virgin Media, TalkTalk and Sky) for a total of 564,462 respondents.

ISPs sampled respondents based on the sampling principles indicated on [pages 7-9 of Ofcom’s Voluntary Codes of Practice \(Residential and Business\) on Better Broadband Speeds Information](#).

This selection represents each broadband product they currently market and must be spread broadly evenly across the footprint in which each service is available (i.e. ISPs serving multiple areas of the UK must not concentrate their tests in a small and unrepresentative number of locations).

The full set of results from the extended embedded data collection is included in the section reviewing technology performance variation across UK regions. A subset of the data was included in the ISP package comparisons, which includes data on packages whose speed tier is not currently measured by SamKnows’ panel or SamKnows’ enabled embedded panel. The included results did not have to be weighted as results from the package group speed tier is only compared to itself or data from a specific package only compared to other packages within the same package group. Data from the embedded data collection was only ever compared to other data from this collection. At no point were direct comparisons made within a package across results from the Whitebox panel and the extended embedded dataset.

The visualisation reviewing technology performance across UK regions shows mean and median, and the upload and download speeds experienced between 8-10 pm across speed tiers. We calculate the summary statistic for each speed tier using data from all the packages in that speed tier. To ensure the statistic represents upload and download speed across the packages in that speed tier in that region, rather than being skewed towards the results of the package with the most data points, we have calculated mean values at the level of individual packages. To create the overall statistic for a given region and speed tier, we have then averaged the mean values for all the packages in a given region and speed tier. For the median values, we have similarly calculated the median for each package individually and then found the median of these values. Below is an example of how this process works.

For example: 140 FTTP 8-10pm download speed in the North East of England region		
	Mean Value	Median Value
Package 1	140.50	139.12
Package 2	143.11	140.16
Package 3	146.25	143.20
Overall data point for region/speed tier	Mean value of those above = $(140.50 + 143.11 + 146.25)/3 = 143.29$	Median value of those above = 140.16

A3. Sources

Sources

[A] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (Less than 10 Mbit/s: 7, above 10Mbit/s and less than 30Mbit/s: 77, above 30Mbit/s and less than 300Mbit/s: 1871, above 300Mbit/s: 600; ADSL: 81; FTTC: 1441; cable: 507; full-fibre: 594). Notes: (1) Data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) Data are collected from multi-thread download speed tests; (3) The above 10Mbit/s and less than 30Mbit/s includes ADSL2+ connections which are not marketed using a connection speed.

[B] Source: Ofcom, based on data provided by the UK's largest broadband providers by retail market share (representing over 90% of the total market), data as at March 2023. Notes: (1) The up-to-30Mbit/s category includes ADSL2+ connections which are not marketed using a connection speed.

[C] Source: SamKnows measurement data for all rurality panel members with a connection in March 2023. Panel base: 2577 (urban ADSL1: 0; rural ADSL1: 7; urban ADS2+: 30; rural ADSL2+: 44; urban FTTC: 964; rural FTTC: 387; urban FTTP: 411; rural FTTP: 183; urban Cable: 499 and rural Cable: 8) Notes: (1) Data have been weighted by take-up rates by speed tier and technology within rurality across FTTC and Cable connections to ensure urban and rural sub-samples are representative of the market by rurality and of the UK overall; (2) Data are collected from multi-thread download speed tests; (3) The bars indicate that there is a 95% probability that the actual average speed for all

corresponding consumers fall within the given range.

[D] Source: SamKnows measurement data for all rurality panel members with a connection in March 2023. Panel base: 2557 (urban: 1925; rural: 632) Notes: (1) Data have been weighted by take-up rates by speed tier and technology across FTTC and Cable connections within rurality to ensure urban and rural sub-samples are representative of the market by rurality and of the UK overall; (2) Data are collected from multi-thread download speed tests.

[E] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (ADSL2+: 74; FTTC 36-38: 458; FTTC 50-55: 102; FTTC 66-67: 787; cable 125: 100; cable 250: 101; cable 350: 100; cable 500: 100; cable 1000: 100, FTTP 67: 77; FTTP 145: 86 and FTTP 300: 60). Notes: (1) All connections, ADSL and FTTC data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) Data are collected from multi-thread download speed tests; (3) Due to the low representation of high-speed cable packages in the UK, broadband provider panel results are used for cable 125Mbit/s, cable 250Mbit/s, cable 350Mbit/s, cable 500Mbit/s, FTTP 67Mbit/s, FTTP 145Mbit/s and FTTP 300Mbit/s.

[F] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (ADSL2+: 74; FTTC 36-38: 458; FTTC 50-55: 102; FTTC 66-67: 787; cable 125: 100; cable 250: 101; cable 350: 100; cable 500: 100; cable 1000: 100, FTTP 67: 77; FTTP 145: 86 and FTTP 300: 60). Notes: (1) All connections, ADSL and FTTC

data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) Data are collected from multi-thread download speed tests; (3) Maximum speed is calculated as the average of the daily maximum speeds achieved throughout the month; (4) Due to the low representation of high-speed cable packages in the UK, broadband provider panel results are used for cable 125Mbit/s, cable 250Mbit/s, cable 350Mbit/s, cable 500Mbit/s, FTTP 67Mbit/s, FTTP 145Mbit/s and FTTP 300Mbit/s. (5) The bars indicate that there is a 95% probability that the actual average speed for all corresponding consumers fall within the given range.

[G] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (Less than 10 Mbit/s: 7, above 10Mbit/s and less than 30Mbit/s: 77, above 30Mbit/s and less than 300Mbit/s: 1871, above 300Mbit/s: 600. Notes: (1) Data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) The above 10Mbit/s and less than 30Mbit/s includes ADSL2+ connections which are not marketed using a connection speed.

[H] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (ADSL2+: 74; FTTC 36-38: 458; FTTC 50-55: 102; FTTC 66-67: 787; cable 125: 100; cable 250: 101; cable 350: 100; cable 500: 100; cable 1000: 100, FTTP 67: 77; FTTP 145: 86 and FTTP 300: 60). Notes: (1) All connections, ADSL and FTTC data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) Data

are collected from multi-thread download speed tests; (3) Due to the low representation of high-speed cable packages in the UK, broadband provider panel results are used for cable 125Mbit/s, cable 250Mbit/s, cable 350Mbit/s, cable 500Mbit/s.

[I] Source: SamKnows measurement data for all national panel members with a connection in March 2023. Panel base: 2557 (ADSL2+: 74; FTTC 36-38: 458; FTTC 50-55: 102; FTTC 66-67: 787; cable 125: 100; cable 250: 101; cable 350: 100; cable 500: 100; cable 1000: 100, FTTP 67: 77; FTTP 145: 86 and FTTP 300: 60). Notes: (1) All connections, ADSL and FTTC data have been weighted by broadband provider package market share, nation and geographic market classification, max attainable speed (FTTC) and distance from exchange (ADSL) to ensure that they are representative of the UK as a whole; (2) Data are collected from multi-thread download speed tests; (3) Due to the low representation of high-speed cable packages in the UK, broadband provider panel results are used for cable 125Mbit/s, cable 250Mbit/s, cable 350Mbit/s, cable 500Mbit/s, FTTP 67Mbit/s, FTTP 145Mbit/s and FTTP 300Mbit/s. (3) The bars indicate that there is a 95% probability that the actual number of disconnections for all corresponding consumers fall within the given range.

[J] Source: SamKnows measurement data for all broadband panel members with a connection in March 2023. Panel base: 9256 ; BT 300Mbit/s full-fibre: 61; BT 145Mbit/s full-fibre: 62; BT 67Mbit/s full-fibre: 58; Sky 145Mbit/s full-fibre: 2136; TalkTalk 150 full-fibre: 3627; TalkTalk 250 full-fibre: 1271; Gigaclear 300 full-fibre: 44; Sky 500 full-fibre: 2148; BT 910 full-fibre: 59; Virgin 1130 cable: 100; Virgin 516Mbit/s cable: 100; Virgin 363Mbit/s cable: 100; Virgin 213Mbit/s cable: 101; Virgin 108Mbit/s cable: 100; Vodafone 67Mbit/s FTTC: 68; TalkTalk 67Mbit/s FTTC: 136; TalkTalk 36Mbit/s FTTC: 90; Plusnet 66Mbit/s FTTC: 12; Plusnet 36Mbit/s FTTC: 87; EE 36Mbit/s FTTC: 90; BT 67Mbit/s FTTC: 198; BT 50Mbit/s FTTC: 100; BT 36Mbit/s FTTC: 99;

ADSL2+: 66. Notes: (1) FTTC data have been normalised to the UK profile by max attainable line speed to ensure that operators can be compared on a like-for-like basis; (2) Data are collected from multi-thread

download speed tests; (3) The bars indicate that there is a 95% probability that the actual average speed for all corresponding consumers fall within the given range.

A4. Glossary of terms

ADSL: Asymmetric digital subscriber line. A digital technology that allows the use of a standard telephone line to provide high speed data communications. Allows higher speeds in one direction (towards the customer) than the other.

ADSL1: The first generation of ADSL, capable of theoretical data speeds of up to 8Mbit/s towards the customer and up to 40kbit/s from the customer.

ADSL2+: An improved version of ADSL, offering higher speeds, especially on shorter telephone lines. In the case of ADSL2+, theoretical speeds of up to 24 Mbit/s can be delivered towards the customer.

Advertised speed: The speed at which broadband services are typically marketed, usually expressed in Mbit/s (megabits per second).

Bandwidth: The maximum amount of data that can be transmitted along a channel.

Broadband: A service or connection generally defined as being 'always on', providing a bandwidth greater than narrowband.

Broadband provider: A company providing broadband internet access.

Broadband speed: The speed at which data are transmitted over a broadband connection, usually measured in Mbit/s (megabits per second).

Cable: Sometimes referred to as hybrid fibre coaxial (HFC) networks, cable networks combine optical fibre and coaxial cable (a cable made up of a conductor and a tubular insulating layer) to carry TV and broadband signals to end users. DOCSIS (data over cable service interface specification) is the technology standard used to deliver high speed broadband over HFC networks.

Contention: A slowdown in performance caused when multiple users share the same

bandwidth within a network and the bandwidth available is less than the aggregate demand.

Download speed: The rate of data transmission from a network operator's access node to a customer, typically measured in Megabits per second (Mbit/s).

DNS: The domain name service (or system) provides a crucial role in the internet. This protocol translates domain names (such as google.com) into the IP addresses that are used to route traffic (e.g. 80.77.246.42). Every broadband provider maintains its own DNS servers through which customers' computers issue queries to translate names into IP addresses. When these servers fail or operate slowly, web browsing and other online activities suffer.

Exchange: The local telephone exchange is the building where all customers' copper telephone lines are connected to enable telephone calls to be switched, and where network equipment is installed which enables customers' data traffic to be routed via an operator's core network to its destination.

FTTC: Fibre-to-the-cabinet. An access network consisting of optical fibre extending from the access node to the street cabinet. The street cabinet is usually located only a few hundred metres from the subscriber premises. The remaining segment of the access network from the cabinet to the customer is usually a copper pair, but another technology such as wireless could be used.

Gbit/s: Gigabits per second. A unit measuring the bit rate. 1 Gbit/s is the equivalent of 1,000 Mbit/s.

G.fast: a technology which allows for superfast internet access over short copper lines which is used to provide some UK fibre to-the-cabinet (FTTC) connections. G.fast can provide speeds greater than those of VDSL.

Headline speed: See 'advertised speed'.

Jitter: The variation in latency. A measure of the stability of an internet connection.

Latency: The time it takes a single packet of data to travel from a user's device to a third-party server and back again. The figure is most commonly measured in milliseconds, and a connection with low latency will feel more responsive for simple tasks like web browsing.

LLU (local loop unbundling): LLU is the process whereby incumbent operators (in the UK these are BT and Kingston Communications) make their local network (the lines that run from customer's premises to the telephone exchange) available to other communications providers. The process requires the competitor to deploy its own equipment in the incumbent's local exchange and to establish a backhaul connection between this equipment and its core network.

Local loop: The access network connection between the customer's premises and the local telephone exchange, usually a loop comprising two copper wires.

Mbit/s: Megabits per second. A unit measuring the bit rate. 1 Mbit/s is the equivalent of 1,000 Kbit/s.

Multi-thread test: A test involving the download of two or more data files

simultaneously – in the case of our research, three files (see Technical Methodology – Annex 1). Multi-thread tests typically record faster speeds than single-thread tests, in particular for higher-speed connections.

Packet loss: The loss of data packages during transmission over an internet connection.

Streaming content: Audio or video files sent in compressed form over the internet and consumed by the user as they arrive. Streaming is different to downloading, where content is saved on the user's hard disk before the user accesses it.

Sync speed: Modem synchronisation speed. The maximum download speed that a line can support according to the way the line is configured by a customer's broadband provider.

Upload speed: Also uplink or upstream speed. Rate of data transmission from a customer's connection to a network operator's access node, typically measured in Mbit/s (megabits per second).

VDSL (Very high bitrate high-speed digital subscriber line): a technology which allows for superfast internet access over copper lines and which is used to provide most UK fibre-to-the-cabinet (FTTC) connection