
High-Level Testing Principles

Annex 5 to the Voluntary Codes of Practice on
Residential and Business Broadband Speeds

1. Introduction

- 1.1 This document is an annex to the Voluntary Codes of Practice on Residential and Business Broadband Speeds (Residential and Business Codes) and should be read alongside them. It outlines the approach that must be used to test the actual speeds received by a sample of an ISP's customers, as outlined in paragraph 2.20 of the Residential and Business Codes. This is to help ensure consistency across Internet Service Providers (ISPs). The proposed technical methodology is broadly consistent with the approach used in our Home Broadband Performance research.¹
- 1.2 The high-level principles in Section 2 are intended to cover all access technologies, as are the requirements of the Residential and Business Codes, which are intended to be technology neutral in outcome. This is on the basis that all methodologies used to measure peak time speed adequately reflect the specific attributes of each access technology.
- 1.3 The more detailed methodology described in Section 3 of this document relates to xDSL, cable and FTTP. Signatories and prospective signatories intending to use other technologies, including fixed wireless or satellite, must discuss appropriate testing methodologies with Ofcom. However, we expect testing methodologies and customer speed estimates for fixed wireless services to follow the codes and the high-level principles set out in the Section 2. The specific factors that lead to speed degradation on this technology would include distance from the mast (including the length of relays to properties outside direct line of sight) and network use in the area.
- 1.4 The high-level principles will be subject to ongoing review.

¹ <https://stakeholders.ofcom.org.uk/market-data-research/telecoms-research/broadband-speeds?a=0>

2. High-Level Principles

- 2.1 The aim of the testing is to determine the effect of contention on actual fixed broadband speeds delivered at peak times and during the quiet hour.

Test principles

- 2.2 ISPs must select a test panel of customers based on the sampling principles outlined in Section 4. This selection must represent each broadband product² 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).
- 2.3 Each panellist must have a unit capable of running the appropriate test software that can measure download and upload speeds received at the customer premises equipment (CPE). The ISP may run additional quality tests³ if they wish to do so, but for the purposes of the code measuring download and upload speeds will be sufficient.
- 2.4 The software must perform daily tests for each panellist during peak time (8-10pm for residential services, 12-2pm for business) and the quiet hour (the time at which the ISP expects the network traffic to be least contended). ISPs must ensure that tests are spread out across each of these periods.
- 2.5 To determine the maximum speed achieved on the panellist's line, ISPs may test throughout the day rather than solely in the expected quiet hour.
- 2.6 The data used to calculate contention must be updated at least quarterly, drawing on the previous three months' speed measurements, although ISPs may update more frequently if they wish to do so.
- 2.7 The download and upload speed 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.
- 2.8 To avoid detriment to panellists, the data used for the tests must not be included in the panellists' data allowance (if any).
- 2.9 To assess the capacity of the user's broadband connection, up to three concurrent transmission control protocol (TCP)⁴ connections must be used.

² Each broadband service with a different advertised headline speed is a different product. Products with the same advertised download speed but a different advertised upload speed or different data allowances must be treated as separate products for the purposes of download speeds if these differences affect downstream contention. In some instances (outlined in paragraphs 3.6-9) it may not be necessary to have a separate panel for each product.

³ E.g. jitter, latency.

⁴ TCP enables two hosts to establish a connection and to exchange data

High-Level Testing Principles

- 2.10 The download and upload tests will consist of downloading and uploading files over a duration of 5 seconds (starting from the end of the ramp-up period⁵ described in paragraphs 2.13 and 2.14). The application layer protocol must be http 1.1.
- 2.11 The files to be used for the testing will be stored on an Ofcom webserver, from where the ISPs can retrieve them for use.
- 2.12 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. We expect ISPs to use their own server for the testing; if there is a need for a centrally managed server, ISPs should discuss this with Ofcom.
- 2.13 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.
- 2.14 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.

⁵ The slow-start part of the congestion control strategy used by many internet connections

3. Calculating speed estimates for different technologies

xDSL Technologies

- 3.1 For each panellist, their mean peak and quiet hour⁶ speeds must be compared against their sync speeds. This comparison must be used to derive, for each broadband product and at a national level, two coefficients representing the mean difference against sync (the national contention coefficient) for the whole product panel: one for peak time and one for the quiet period. See the table below for further detail.

Cable/ FTTP Technologies

- 3.2 The mean peak and quiet hour speeds⁷ for each panellist must be collated at a relevant network level to create a data distribution for each group at peak time and quiet hour. The relevant aggregation level for these technologies is currently nationwide.
- 3.3 To calculate the estimated speeds for a customer, the measurements must be used as set out in the following table. This approach must be used for both upload and download speeds:

Metric	Format	xDSL	Cable/ FTTP
Normally available	Range	<p>The 20th and 80th percentiles of the sync speed of lines characteristically similar to the customer's line and on the relevant product are derived.</p> <p>The national contention coefficient for the relevant product from peak time is applied to these speeds.</p> <p>The resulting range (of sync speeds adjusted to account for contention at peak time) is presented to the customer.</p>	<p>The 20th and 80th percentiles of the peak-time speed of lines on the relevant product and at the network level relevant to the customer are derived.</p> <p>The resulting range (of actual speeds received by similar customers at peak time) is presented to the customer.</p>

⁶ Or, if testing throughout the day in accordance with paragraph 2.5, the mean of the highest recorded speed from each day

⁷ Or, if testing throughout the day in accordance with paragraph 2.5, the mean of the highest recorded speed from each day

Minimum	Point	<p>The 10th percentile of the sync speed of similar lines on the relevant product is derived.</p> <p>The national contention coefficient for the relevant product from peak time is applied to this speed.</p> <p>The resulting figure is presented to the customer.</p>	<p>The minimum speed is at least 50% of the headline advertised speed for the product.⁸</p> <p>The resulting figure is presented to the customer.</p>
Maximum (quiet hour maximum)	Point	<p>The 80th percentile of the sync speed of similar lines on the relevant product is derived.</p> <p>The national contention coefficient for the relevant product from the quiet hour is applied to this speed.</p> <p>The resulting figure is presented to the customer.</p>	<p>The 80th percentile of the quiet-hour speed of lines on the relevant product and at the network level relevant to the customer is derived.</p> <p>The resulting figure is presented to the customer.</p>

Examples of how to apply the methodology to obtain download speed estimates

76Mbit/s FTTC package

A representative sample of the product’s customers is tested. For each customer, their mean peak time speed is compared to their current sync speed.

For the whole panel, the mean difference between peak time and sync speeds is calculated to be a 20% reduction. The difference between quiet hour and sync speeds is calculated to be a 10% reduction.

A new customer provides their address from which the relevant group of similar lines is identified. The network infrastructure provider tells the ISP that the 80th, 20th and 10th percentiles of this group’s sync speed distribution are 60Mbit/s, 55Mbit/s and 40Mbit/s respectively.

To derive their speed estimates, the contention measure is applied to these speeds as follows:

Normally available speed (20th to 80th percentiles, adjusted for peak):
 55Mbit/s to 60Mbit/s, reduced by 20% = **44Mbit/s to 48Mbit/s**

Minimum guaranteed speed (10th percentile, adjusted for peak)

⁸ If the product does not have an advertised speed, this should be based on the ISP’s internal specifications for the service

40Mbit/s, reduced by 20% = **32Mbit/s**

Maximum speed (80th percentile, adjusted for quiet hour)

60Mbps, reduced by 10% = **54Mbit/s**

200Mbit/s cable package

A representative sample of a product's customers is tested nationally. Two data distributions are created: one with each panellist's mean peak time speed and one with each panellist's mean quiet hour speed.

To derive the new customer's speed estimates, this data is used as follows:

Normally available speed

The 20th and 80th percentiles of peak speeds are derived

e.g. 160Mbit/s to 185Mbps

Minimum guaranteed speed

At least 50% of the 200Mbit/s product speed

e.g. 100Mbit/s

Maximum speed

The 80th percentile of quiet hour speeds is derived

e.g. 205Mbit/s

4. Sampling

- 4.1 ISPs must test speeds that represent their typical customer speeds, using a sample large enough to be statistically meaningful, with a geographic spread broadly representative of the exchanges (or equivalent) served by the ISP, covering small and large exchanges (if applicable), across a range of line speeds (if applicable), and across representative VLANs (if applicable). This is to ensure that the speed data captured is broadly representative of the network.
- 4.2 ISPs must test the speeds received by a representative sample of their customers for each broadband product. A separate sample must be created for each package that an operator currently sells to new customers.
- 4.3 The sample used for each package must be large enough to be statistically significant at a 95% confidence interval, with a precision of $\pm 5\%$. Where an ISP is unable to gather a testing panel large enough for this degree of precision, they must use the lower bound of the data range.

Example sample sizes

<i>Size of CP network</i>	<i>Sample coverage by exchange (or equivalent)</i>
Geographically limited, fewer than 50 exchanges	All exchanges tested (1-2 lines at each) will give a census result
Geographically limited, fewer than 200 exchanges	50 exchanges (1-2 lines at each) will give a wide estimate
Geographically limited, more than 200 exchanges	100 exchanges (1-2 lines at each) will give a wide estimate
Nationwide, large number of lines across many exchanges	500 exchanges tested (ideally 3 lines at each) will give a suitably precise estimate
Nationwide, small number of lines across many exchanges	<i>Likely to fall under 20,000 threshold noted in 3.6/9</i>

Testing level

- 4.4 Testing must be sampled at an appropriate network level that reflects differences in how contention affects customers. For xDSL, cable and FTTP technologies this is currently a national level, but this may be different for other access technologies.

New/small packages

- 4.5 ISPs selling new packages with small customer bases (generally fewer than 20,000 subscribers) may use contention measures from a similar existing product, if available and appropriate.

4.6 If not available, suitable alternative approaches must be discussed with Ofcom.

New/small ISPS

4.7 Smaller ISPs that are subsidiaries of larger providers, may use contention measures from that provider's similar products, if they are available and appropriate.

4.8 Otherwise, ISPs with small customer bases (generally fewer than 20,000 per product) or narrow geographic coverage must discuss suitable alternative approaches with Ofcom. This may include sampling by technology, rather than by product.⁹

⁹ E.g. an ISP with an ADSL20, FTTC40 and FTTC80 product range would need one panel for ADSL and one for FTTC, and would use the overall FTTC contention metric for the 40 and the 80 package.