

Net Neutrality Regulatory Assessment Methodology

5 October, 2017

Table of Contents

1. Executive summary	2
2. Introduction	2
3. Measuring Internet access service quality	4
3.1 IAS speed measurements	5
3.2 Delay and delay variation measurements	9
3.3 Packet loss measurements	10
4. Detecting traffic management practices that impact individual applications	10
4.1 Connectivity measurements	10
4.2 Detecting practices that impact QoS of individual applications	12
5. End user dependent factors that may impact the measurement results	14
5.1 End user initiated measurements	14
5.2 End user environment	15
5.3 Hardware and software information retrieval methods	16
5.4 Measurements data filtering	17
6. Measurement results assessment	17
6.1 Data validation	18
6.2 Speed assessment for end users	18
6.3 Market level aggregation	20
6.4 Individual applications using IAS	21
7. Certified monitoring mechanism	22
7.1 Guidance on criteria regarding certified monitoring mechanism	23
8. References	23

1. Executive summary

This document contains BEREC regulatory assessment methodology in order to provide guidance to National Regulatory Authorities (NRAs) with the monitoring and supervision of the net neutrality provisions of the Regulation 2015/2120 [1], implementation of net neutrality measurement tools by NRAs on an optional basis, and contribute to standardization of net neutrality measurement methodology. The work builds upon previous BEREC guidance on net neutrality, Internet access service (IAS) quality monitoring and best practices.

Chapter 3 provides guidance on a harmonised quality of service measurement methodology. It is targeted to maximising measurement accuracy and to enable the comparison of measurement results between different member states. The speed measurement is based on multiple HTTP connections and the document describes the calculation of the speed. The document also illustrates the difference between the two options of calculating speed based on TCP or IP payload. This document also defines measurement metrics for delay, delay variation and packet loss.

Chapter 4 gives recommendations on methods for detecting traffic management practices that impact individual applications and suggests other indicators of performance closer to the user experience. It includes recommendations for detecting traffic management practices that affect the connectivity and ultimately a possibility to use and provide individual applications. The document describes also recommendations for detecting traffic management practices that affect the quality of individual applications like the prioritisation and/or throttling of specific applications.

Chapter 5 describes the most important factors that should be taken into account when assessing the measurement results and gives guidance on information collection. Thus, a number of end user environment factors may impact the results. These factors include for example Wi-Fi usage, modem and computer performance and radio conditions when measuring speed for mobile subscription.

Chapter 6 provides recommendations for validation of the collected measurement results. It also provides some further guidance on how the speed measurement results should be assessed in comparison to the contractual speed values for end users. The topic of data aggregation for market level assessment purposes is discussed and guidance on monitoring the general IAS quality (IAS as a whole and effect of specialised services on IAS) as well as individual applications using IAS is provided.

Finally Chapter 7 gives guidance on the criteria that NRAs could take into account when providing their own certified mechanism or certifying a third party mechanism.

2. Introduction

BEREC has developed this regulatory assessment methodology in order to help NRAs in the monitoring and supervision of the net neutrality provisions of the Regulation [1] based on various net neutrality measurement tools and harmonised measurement methodology for quality of service indicators. Other aspects of harmonisation such as sampling and validation of collected measurements have not been fully considered here.

Furthermore, it is also foreseen that this BEREC net neutrality measurement methodology can contribute to the work of standardization bodies. The methodology builds upon previous BEREC guidance on net neutrality, internet access service (IAS) quality monitoring and best practices [2].

Under the Regulation, NRAs may have several objectives in measuring IAS:

- Measurement tools can be used for detecting traffic management practices which may or may not be allowed (art. 3(3) of the Regulation).
- Measurement tools can be used for the establishment of what the 'general quality of IAS' is. This is relevant to the assessment of whether services other than IAS (in the meaning of article 3(5) of the Regulation) can be provided.
- Measurement tools may be part of a monitoring mechanism certified by the NRA as referred to in article 4(4) of the Regulation.

NRAs may have additional objectives in measuring or detecting certain practises related to the IAS. The objectives could for example be defined based on end-user reporting. BEREC notes that it is for NRAs to determine the most appropriate measurement tools to serve their objectives. Different objectives may lead to the use of different measurement tools.

In this document, BEREC describes the methodology for the measurement of IAS speed to enable NRAs to assess IAS performance compared to the contractual minimum, normally available and maximum speed values. The methodology also gives guidance on some criteria that NRAs could take into account when providing their own measurement tools as a certified mechanism or certifying a third party mechanism in accordance with the Regulation and BEREC NN guidelines.

This document aims to describe a measurement methodology that could be combined with a crowdsourcing approach so that it would be possible to provide measurement tools for end users. For in-browser or app based crowdsourcing measurement tools it is hard or even impossible to have full control over all the factors such as the end user environment that impact measurement results. This introduces a possibility for error in measurement results that cannot be fully avoided. This methodology provides guidance on how to increase the accuracy and reliability of such measurement results. This is discussed in Chapter 5.

BEREC recognises the standardised measurement approaches from ETSI, ITU and IETF, however BEREC places the greatest emphasise on the practical implementation of the measurement methodology in a crowdsourced scenario, where measurements can be run by any end-user and the metric being measured reflects as closely as possible real end-user experience of Internet usage.

As proposed in the 2012 BEREC NN QoS Guidelines [2], the measurement methods shall encompass both the IAS as a whole as well as individual applications delivered via IAS. The methodology supports both IPv4 and IPv6 - this topic is further discussed where necessary.

3. Measuring Internet access service quality

The aim of this chapter is to specify the best practice methodology based on the combined goal of maximising measurement accuracy balanced against the need to be able to facilitate easy access to the measurement tool for the general public ensuring that the measurement results are comparable between different member states.

Results of these measurements can be also used for the following purposes:

- Empowering the end user to validate the commitments made to them from their IAS provider.
- Monitoring the general IAS quality and confirming that the performance of IAS is developing sufficiently over time when taking into account technological evolution.
- To support the detection of traffic prioritisation and/or throttling of selected applications compared to other applications running over IAS.
- NRAs may also use the data to increase transparency (e.g. interactive maps showing performance in a geographic area).

According to BEREC NN guidelines [3] paragraph 166, "Measurements should be performed beyond the ISP leg." and speed should be calculated "based on IP packet payload, e.g. using TCP as transport layer protocol". According to the NN guidelines paragraph 140, "Speeds should be specified on the basis of the IP packet payload or transport layer protocol payload, and not based on a lower layer protocol".

This methodology is targeted to measure IAS quality in both the upload and download direction. It is worth noting that IAS speed is just one component of the performance experienced by the end users, since different applications have different protocol overheads and different requirements related to IAS delay, delay variation and packet loss.

For both measurement tasks - IAS as a whole and individual applications using IAS - the fundamental precondition is that measurements are performed at the edge of the network which provides the IAS (i.e. end user premises for fixed access or via the radio access for Mobile IAS). When initiating a connection, mobile networks need some time to allocate the resources. This delay may impact the measurement results and therefore the measurement results from the first few seconds may need to be discarded.

Where measurements are performed against a test server, this server should be located outside the IAS network¹. It should have adequate connectivity between the server and the IAS provider to avoid influencing measurements. Typically this can be achieved by locating the measurement server at the national Internet exchange point (IXP). Depending on the specific national situation, measurement servers may be located at more than one IXP. There might also be specific reasons for its placement elsewhere. The hardware running the measurement server is connected as close to the IXP switch as possible. This means that the number of physical paths between the main IXP switch and the gateway serving the system should be kept at a minimum. This is applicable whether the implementation runs inside the

¹ This is in line with the architectural aspects of a practical measurement system as described by BEREC in a previous publication [5].

network of a hosting provider or directly on hardware under control of the NRA itself, in a separate AS. The goal is to minimize latency added due to the communication paths.

BEREC recommends that NRAs keep in mind that traffic to the server is not handled differently to other internet traffic.

In cases where both the server and client have visibility of test results (e.g. speed measurements), in general the receiver of the data is considered to be the authoritative source of measurement results, but it is recommended that both the client and server measurements are stored for analysis.

Monitoring mechanisms should mitigate, to the extent possible, confounding factors which are internal to the user environment. Examples of these factors include existing cross-traffic and the usage of Wi-Fi based interfaces. This topic is discussed separately in chapter 5.

The assessment of measurement results is discussed further in chapter 6 and the certified monitoring mechanism is further discussed in chapter 7.

3.1 IAS speed measurements

3.1.1 Speed measurement overall methodology

In order to maximise compatibility in a real world environment, it is recommended to measure upload/download speeds based on the time to execute a set of controlled data transfers over TCP/HTTP. In this way, the speed can be measured based on transport layer protocol payload as referred to in the BEREC NN guidelines [3] paragraph 140.

This methodology is supported by the broadest range of platforms, and can be implemented within a web browser or within the restricted sandbox of an on-device app. As such it is considered to be the adequate balance between the competing demands of accuracy, platform agnosticism, ease of implementation and transparency.

These requirements and the experience from the existing measurement systems lead to the choice to use TCP/HTTP as the measurement protocol. The decision is also supported by the widespread use of TCP by the normal internet applications and services and therefore it reflects the typical IAS usage by the end-users.

The reason to recommend the use of HTTP (1.1 or 2.0) as the TCP payload is to mitigate any connectivity restrictions (such as might be introduced by intermediate firewall or proxy servers) which could result from the choice of a less commonly used protocol/port.

The use of Transport Layer Security in the form of HTTPS is optional, but provides the advantage of preventing any manipulation from intermediate proxy servers².

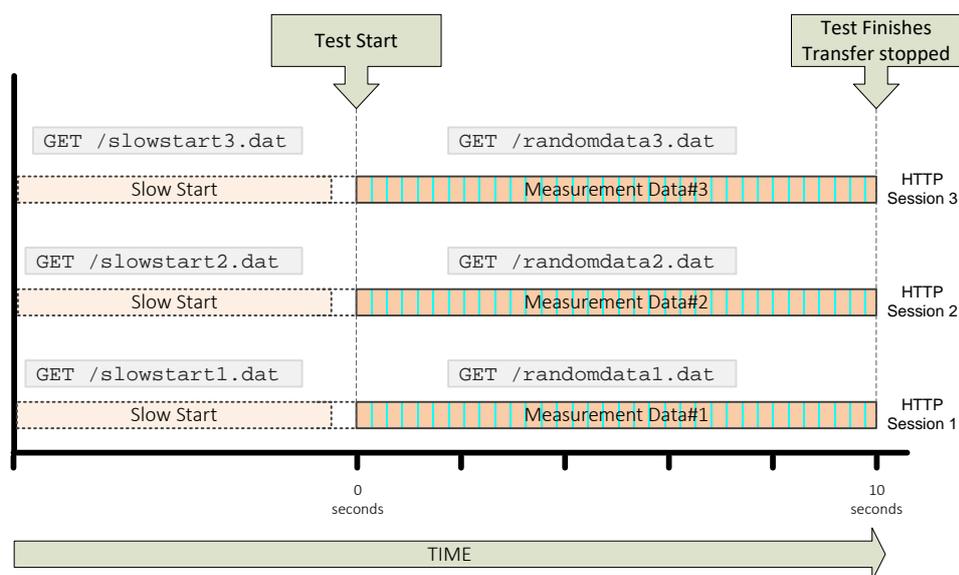
In order to saturate the path, it is recommended to use at least 3 persistent HTTP connections. When developing a measurement tool, it is recommended to verify what setup is needed to measure multigigabit IAS subscriptions. These connections should all have completed the

² Proxies might be used in office environments.

TCP slow start phase before measurement commencement to maximise throughput and ensure that the measurement is as representative as possible. The test should be stopped after a pre-defined interval and the transfer speed calculated by the recipient³ based on the data transferred during that interval, specifically the duration between the arrival of the first packet of the test to the arrival of the last packet, from the perspective of the recipient.

BEREC recognises that packet loss and packet retransmission has a negative impact on the throughput of each TCP connection and hence the IAS speed.

The following diagram illustrates an example implementation of a download test based on 3 HTTP connections.



The following points should be noted in relation to the diagram above:

- 3 persistent HTTP connections are started at the same time;
- To mitigate the effect of TCP slow-start, an initial retrieval of a slowstartX.dat file is made to maximise subsequent throughput;
- Once the transfer of slowstartX.dat file is finished for each connection, the real test commences (using same TCP socket) without delay;
- All .dat files referred to above contain random data, which cannot be compressed and
- The test is stopped after a total of 10 seconds, and the valid measured upload/download volume is based on the total transferred volumes in Measurement Data#1, Measurement Data#2 & Measurement Data#3.
- The test duration, filenames used and number of connections are not proscriptive and this diagram depicts an example for illustration purposes

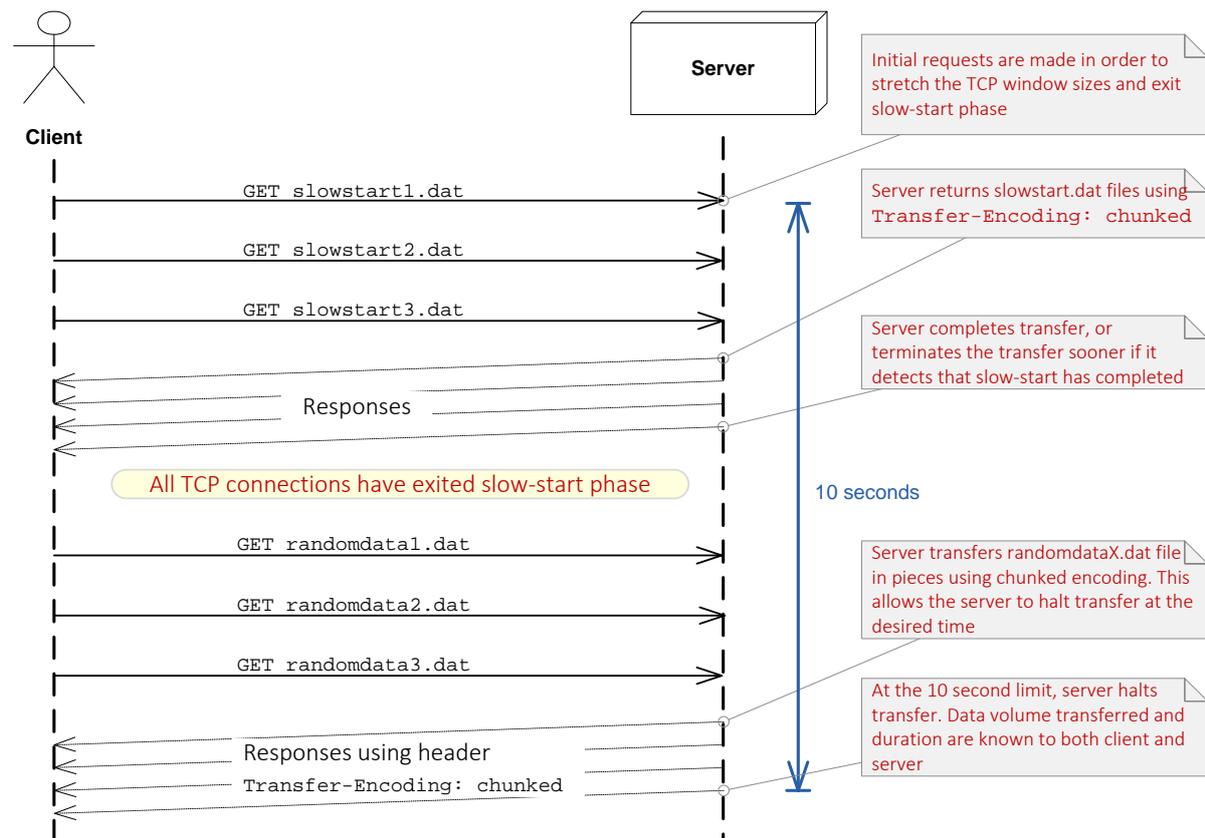
It is recommended that the HTTP transfers are made using chunked transfer encoding which enables the sending side to send chunks of data of arbitrary size and stop the transfer at the appropriate time. For example the sender could send the measurement data in fixed 1MB

³ For upload tests the recipient is the measurement server, whereas in download tests the recipient is the client (end user).

chunks, stopping when the test duration has fully passed; or a more sophisticated scenario could be employed where a variable chunk size is determined by the sender based on the testing time remaining and transfer speed observed.

Irrespective of the implementation chosen, it is important that each chunk is served immediately and without any delay, and any additional data volume transferred as a result of Chunked Transfer Coding (or chunk extensions) is included in the speed calculation.

The diagram below shows the HTTP transfers in more detail in a message sequence chart format.



The following sections discuss the ways to calculate the actual data transferred under this methodology.

3.1.2 Calculating speed based on TCP payload

Calculating the TCP payload is relatively straightforward as compared to calculating the IP Payload (see section 3.1.3). For a given HTTP connection, both the client and the server are mutually aware of the data volume transferred. This data volume will vary for each connection due to the recommendation that the measurement test duration is fixed.

Note that the amount of data transferred will also include the HTTP headers and any Chunked Transfer overhead, both of which must be taken into account and included in the speed calculation. It is recommended that in the unlikely event that the exact size of these headers is not known, that a fixed 500 byte value is added to the total file size as an approximation.

The error introduced by this approximation is considered to be negligible except in cases where the test is run very briefly or on extremely slow links.

3.1.3 Calculating speed based on IP packet payload

Calculating speed based on IP packet payload is more complex due to the fact that most platforms don't allow clients to access this information directly, so it must be calculated based on certain assumptions, and the results of this calculation is an approximation.

Since the measurement client can only be guaranteed to know the HTTP payload volume (i.e. the size of the file transferred), it would be necessary to calculate the number of packets required to transfer this TCP payload and then use this number to calculate the volume of TCP headers.

However the number of packets is a function of the TCP Maximum Segment Size (MSS), which is itself a function of the Maximum Transmission Unit size (MTU).

In addition to the above, the potential presence of TCP options, both during TCP handshake and in established connections⁴, introduces the possibility that the TCP header size is not fixed which further complicates the calculation.

The result of these factors is that it is impossible to accurately calculate the IP Payload volume from the TCP Payload volume. Therefore an adequate safety margin should be taken into account. Example calculations for the overhead are shown in the following tables for various sample values of MTU and TCP header size for both IPv4 & IPv6.

IPv4 (no IP options)	MTU	
	1500 Octets	1280 Octets
Average TCP Header Size	1.37%	1.61%
20 Octets (no TCP options)	1.37%	1.61%
40 Octets (Average 50% of max)	2.78%	3.28%
60 Octets (Max TCP options)	4.23%	5.00%

IPv6 (no IP options)	MTU	
	1500 Octets	1280 Octets
Average TCP Header Size	1.39%	1.64%
20 Octets (no TCP options)	1.39%	1.64%
40 Octets (Average 50% of max)	2.82%	3.34%
60 Octets (Max TCP options)	4.29%	5.09%

Note that these tables are intended to provide an illustration of the potential impact of these variables; however it's expected that in practise the MTU will generally be very close to 1500 octets and the average TCP header size will be close to 20 bytes. Only the two percentage values highlighted in each table are considered relevant in most real world cases.

⁴ e.g. TCP Selective Acknowledgment or Timestamp options

Therefore 3% TCP header overhead can be considered to include an adequate safety margin and it can be used in calculating the IP packet payload. The IP packet payload calculation is done by adding this 3 % value to the speed calculated based on TCP packet payload.

However it should be noted that BEREC considers that TCP payload volume is the most reliable one to use when calculating the upload/download speed.

3.1.4 Miscellaneous Details

It should be possible to run measurements both over IPv4 and IPv6.

Both download and upload speeds should be measured in the same manner and reported in bits/second (e.g. kbit/s or Mbit/s). Note that conversion factors between mega and kilo shall be base-10 rather than base-2 (i.e. 1 Mbit/s = 1000 kbit/s rather than 1024 kbit/s)

TCP/HTTP characteristics and options

Where possible, it is recommended to mitigate the effect of the following inherent HTTP and TCP characteristics which could otherwise introduce error to speed measurements.

TCP connection speed limit: As the bandwidth of an individual TCP connection is limited to the bandwidth-delay product of the path in question, it is necessary to utilise multiple TCP connections in order to saturate that path.

HTTP considerations: Generally the use of HTTPS is recommended. When using plain HTTP, either the appropriate HTTP headers to prevent caching should be used, or unique URIs should be used.

3.2 Delay and delay variation measurements

In principle, any kind of IP packet could be used for latency measurements (e.g. ICMP, UDP or TCP). However the following considerations should be taken into consideration:

- Operating systems normally require administrator (root) privileges for sending ICMP packets. Also ICMP packets can be blocked by firewalls and antivirus software and hence they cannot be relied upon.
- TCP packets (after connection setup) are subject to flow control. The timing of the TCP session establishment 3-way handshake could be used to measure the delay however this is not routinely measured by the TCP stack and made available to applications.
- In a web browser environment it is difficult or even impossible to send/receive UDP packets.

It is recommended that delay is measured using:

- UDP with ICMP or TCP as fall back option,
- at least 10 measurements, and
- calculated as an average of recorded round-trip time values (typically expressed in milliseconds).

The measurement server should return any UDP packet payload immediately, allowing the client to calculate delay. The Unix echo service could be used for this function.

The measurement setup should be insensitive to (user) clock changes during the measurement.

It is also recommended that the delay variation (jitter) is calculated as mean deviation based on the samples collected for the delay measurement.

Calculation shall be based on the algorithms used in the Linux ping utility which is based on 4.3BSD

For example:

```
--- qostest.eu ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9014ms
rtt min/avg/max/mdev = 5.317/5.442/5.727/0.121 ms
```

3.3 Packet loss measurements

If a packet is not received back within a certain timeout (e.g. 3 seconds), it is considered as lost for the purpose of packet loss measurements.

Due to the effect of packet loss on TCP connections, even the low level of packet loss observed in modern networks⁵ can result in significant performance degradation. Therefore it is evident that 10 or even 100 measurements may not yield packet loss accurately. It is therefore recommended to send a large number of IP packets (e.g. at least 1000).

Delay and packet loss measurements are typically performed over a longer period of time in order to allow for the time varying nature of network performance in packet-switched networks.

However, the principle of running measurements of long duration conflicts with the crowdsourced user-initiated measurement concept. End users will not accept an extended waiting time for the presentation of results.

Conversely, short duration tests can only provide an indication of whether the measurement was done during stable network conditions; whereas long measurement intervals are preferred for the meaningful measurement of the stability of the IAS.

While long duration tests are preferred for delay and packet loss measurements, this likely introduces the need for a measurement client running permanently in the background.

4. Detecting traffic management practices that impact individual applications

4.1 Connectivity measurements

This chapter describes recommendations for detecting traffic management practices that affect the connectivity and reachability of individual applications. This document focuses on

⁵ In modern FTTH networks packet loss may vary e.g. between 0.001 % and 0.21%.
Source: Performance Within A Fiber-To-The-Home Network,
<http://www.sigcomm.org/sites/default/files/ccr/papers/2014/July/0000000-0000002.pdf>

measurement methodology and does not give indications which traffic management practices are allowed and which are not. This use case includes the detection of blocked applications and content (e.g. network based content filtering, such as ad-blocking and blocked web content). In addition end users could be restricted from using and providing applications by blocking communication ports, URLs and IP addresses. Many applications can be blocked by blocking the communication port used by the application. Therefore the connectivity measurements described below are an essential part of the net neutrality assessment methodology and should be used according to the need in each national market.

4.1.1 Blocked ports

Blocked ports can be detected by establishing a connection to the port being tested, using the transport protocol in question. With TCP, a port can normally be considered as being open if communication is possible. Given the connectionless nature of UDP, a measurement system must define a feedback mechanism that tells whether the packet was received.

Measurement tools should be able to test for blocked ports at least over the following:

- IPv4 and IPv6;
- TCP and UDP;
- Uplink (connections from the end user to internet host) and downlink (connections from the internet towards the end user); and
- Any UDP or TCP port number.

It is also worth noting that network address translation (NAT) which might be used by ISPs affects downlink connectivity such that by default all communication ports are blocked in the downlink direction. This must be taken into account when assessing the measurement results.

It is important to take into account that the end user environment (especially firewalls) may also affect the results. However, in the case of crowdsourcing approach it may be possible to compare large number of results from different end users. Setups or disturbances in the end user environment may produce measurement results that incorrectly indicate certain traffic management practices. In case a large number of measurements indicate the same traffic practice, the likelihood that these practices are indeed occurring due to the operator's network setup increases.

4.1.2 IP addresses blocking

The measurement tool must be able to perform this test using both IPv4 and IPv6 protocols. The purpose of this test is to detect if certain IP addresses are blocked. This check is executed by attempting to connect to well-known ports on the address being tested against. Therefore it should first be known that the tested IP address answers requests on a certain port number and the test would most used to check complains. The test methodology is similar to the port-oriented connectivity check described in the previous chapter, but in this case the focus is on the specific IP address.

A successful connection to any port (or indeed any response from that address) is not sufficient to detect that the IP address is not blocked, since some ISPs could use middle-boxes to simulate a connection, and even answer on the established connection. Therefore, it is recommended to also send some data and verify the integrity of the received data.

If the connection cannot be established or the received data is not as expected, a new measurement should be performed using a VPN to access a proxy server outside ISPs control so that the ISP does not see the real destination address. If the connection via proxy is successful, this can be seen as an indication that something in the ISP network is blocking the IP address.

4.1.3 DNS manipulation

DNS manipulation refers to a situation where a DNS reply is received (on an A or AAAA request) which falsely indicates that the domain is unknown or where an incorrect IP address is returned. The result of this manipulation is that the client is redirected to a different address.

DNS manipulation can be detected by analysing the responses to DNS requests on known targets (e.g. DNS records of specific domains under the control of the NRA).

Note that end user environment (especially firewalls) may affect the results. However, in the case of a crowdsourcing approach it may be possible to compare thousands of results from different end users and using different DNS resolvers which could solve the problem.

4.1.4 Detection of an HTTP proxy

An HTTP proxy is a middlebox that is inserted into the path for end users' HTTP connections, which may be used to filter or modify traffic. A HTTP proxy can be transparent or otherwise hidden.

A transparent proxy is a middlebox deployed by the IAS provider which acts as an intermediary between the client and the target web server. Typically, the ISP routes HTTP traffic via the proxy without user action or knowledge.

A transparent HTTP proxy might be detected by checking the HTTP headers for proxy specific content (HTTP_VIA, VIA, FORWARDED, CLIENT-IP...).

The HTTP (TRACE) request headers could also be checked for modification between the client and the server and if the intercepting proxy does a DNS lookup on a fake host header.⁶ A hidden proxy could be detected by a cache test⁷.

Some HTTP proxies can be detected by connecting to a target domain and checking that the web resource is available and verifying that the content is identical to the content received over a proxy outside the ISPs control⁸.

Finally, it may be possible to detect a HTTP proxy by inspecting properties of the sent traffic (e.g. TTL-flag of the IP packet).

4.2 Detecting practices that impact QoS of individual applications

The purpose of these measurements is both to suggest other indicators of performance closer to the user experience, which could give to the consumer some easily understandable criteria to help him to take enlightened decisions, and to detect the prioritisation and/or throttling of

⁶ ProxyDetect: <https://github.com/cyberis/ProxyDetect/blob/master/proxydetect.pl>

⁷ For an example how this could be done, see: www.lagado.com/tools/cache-test

⁸ For an example how this could be done, see: <https://ooni.torproject.org/nettest/http-requests/>

specific applications. These traffic management practices may be detected by measuring some of the KPIs described below and comparing the results based on the following variations:

- Comparison of the same KPIs related to similar applications for the same IAS subscription,
- Comparison of the KPIs for the same application using an equivalent subscription from another IAS provider, and/or
- Comparison of the KPIs for the same application and the same IAS subscription but using a VPN (see section 6.4).

These measurements can be performed on a regular basis for selected applications, websites or platforms or in targeted situations as needed.

4.2.1 *Web browsing*

In order to assess web browsing quality of service, the time (in seconds) needed to load a web page for the first time could be a good indicator. A first loading experience can be simulated by the measurement tools e.g. by forcing local cache clearance or to configure the web server appropriately.

One option can be to use a normalised reference page (e.g. ETSI (mobile) Kepler page)⁹ from a dedicated web server. This avoids bias by fixing the page size and associated page elements, and removes the dependency on the performance of a web server on the open internet. A second approach consists of measuring the time taken to reach the HTML and referenced resources from a page of a real website. The panel of chosen websites can include popular websites, government website, etc.

A significant increase of the time needed to load the web page can be an indication of throttling or de-prioritisation.

4.2.2 *Video streaming*

In order to assess video streaming QoS, two options could be considered:

- To simulate a data stream comparable to a normal video streaming session on which standard measurements of IAS quality can be carried out (bitrate, latency, etc.).
- To launch a video streaming session on an existing public streaming platform.

Several characteristics related to the video streaming quality can be measured:

- The negotiating time used for exchange of setup information for the streaming session between the request and the beginning of the flow;
- The number of breaks - number of cuts that require a renegotiation; and
- The duration of the cuts - accumulated durations of the cuts that occurred during one streaming session.
- In the case of live streaming quality, the delay introduced by encoding the live stream.

It is also possible to measure video streaming QoS by measuring the video streaming bitrate instead of measuring impacts such as decreased video quality, increased negotiation time and

⁹ ETSI Kepler, <https://docbox.etsi.org/STQ%2FOpen%2FKepler>

the number and duration of cuts. This bitrate can be compared to the IAS speed measurement results, while taking into account that bitrate can also be influenced by player capabilities and requirements. Thus the use of adaptive bitrate in video streaming in some terminals may show lower bitrates, regardless of the performance of the network itself.

One possible implementation may consist in launching a video that is longer than a defined minimum length (e.g. 20 seconds) with an encoded bitrate set to the maximum available value. If there is a break, the test is repeated with a lower bitrate. The test is repeated at decreasing bitrates until there is no break during the playback. The highest bit rate at which the playback is successful is then considered to be an indicator of video streaming quality of service. Such a test should be repeated to get more accurate results.

By using a real video streaming platform, it can be difficult to differentiate between video site performance and ISP performance as some video sites might use different streaming servers for different ISPs. Nonetheless, the simulation of a transport-channel is less representative to the end-user experience and may be difficult to implement.

4.2.3 Other use cases

Other use cases could be the subject of future study, following similar methodologies to those described previously. Such candidate use cases include Voice over IP (VoIP), audio streaming or peer-to-peer file sharing and any other future applications not yet released.

5. End user dependent factors that may impact the measurement results

This chapter describes recommendations for NRAs on identifying the end user environment factors which may affect measurement results and in some cases, minimising impacts to the measurement results.

5.1 End user initiated measurements

The main challenge for this type of measurement is the potential impact of the end user environment on the measurement results. Some end users might not be aware of the potentially negative influence of their own terminal equipment and home network but assume that all issues are on the access service or content provider's side.

Based on this, it is important that end users are informed of how to properly set up their terminal equipment to minimise error. It is also crucial to inform the end user that conducting multiple QoS measurements will provide a more representative view of IAS performance, and better analysis and assessment of measurement results.

Moreover, in order to ensure that measurement results are accurate, a measurement server must protect itself against overload, to prevent a case where too many simultaneous measurement sessions cause interference. Therefore any measurement server must implement access control and queueing such that excessive measurement requests are responded to with a temporary delay, resulting in clients waiting until the measurement server is ready.

5.2 End user environment

The end user environment consists of many elements, some of which could limit IAS performance. These limiting factors are listed and described below in two separate subsections differentiating between fixed and mobile environments.

It is recommended that, when available, the measurement client retrieves the required data from the local hardware and operating system (computer and modem router) and from the IAS provider both before and while running measurements.

A prerequisite of such data collection is the informed consent of the user, which should be required by the measurement client before measurement commencement.

5.2.1 Fixed environment

The following sections outline the main issues which might prevent an accurate performance measurement, so these issues should be taken into account when assessing measurement results in fixed environments.

Performance of the modem router

If the performance of the modem and home network used to connect to the internet is not able to deliver at least the contractual bandwidth, the measured performance may not accurately reflect the IAS performance. It should be noted that in some cases the modem is not provided by the IAS provider, so any performance gap may not be their responsibility.

Type of the link

If the measurements are not carried out through a wired connection (via the model/router's Ethernet port) but through another link type that might add extra delay, packet loss or bandwidth reduction (e.g. Wi-Fi, powerline or wireless repeater), the measured performance may not correspond to the IAS performance.

Performance of the computer (CPU and RAM load)

If the load upon the computer in terms of RAM and/or CPU utilisation is too high, the measured performance may not correspond to the IAS performance. This may happen when certain software or applications are not closed down before starting the measurements.

Version of the computer operating system

Outdated operating systems might not include the latest performance tuning patches, and the increased likelihood of automatic updates being downloaded could slow down the transmission rate.

Simultaneous usage of other software like antivirus and firewalls

If background software like virtual private network (VPN), anti-virus, content based filtering (e.g. parental control), firewall and/or any local DNS manipulation is active when running the measurement tasks, the measurement results may not correspond the IAS performance. This is particularly important in detecting traffic management practices that impact individual applications.

Cross traffic

If cross traffic generated in parallel with the traffic of the measurement client, such as download/upload of data, music streaming, IPTV and videoconferencing, then the measured performance may not correspond the IAS performance, and this should be taken into account when assessing the measurement results. Also traffic generated by software updates running on background may impact the measurement results.

The cross traffic may be generated by the same computer that is running the measurement client or other network nodes.

5.2.2 Mobile environment

As above, the following sections outline the main issues which might prevent an accurate performance measurement, so these issues should be taken into account when assessing measurement results.

Performance of handset model

The performance of the handset model involved in the measurements can affect the measurement quality. Different devices perform differently. Therefore care should be taken not to blend data from different devices in mobile environments.

The radio connection quality

The available speed depends on the quality of the radio conditions. Therefore it is important to retrieve and store the information on the radio conditions which prevailed during the measurement. The available radio parameters vary between different mobile network technologies and operating systems. Therefore it is recommended to retrieve available parameters, provided by the mobile handset, e.g.: RSSI, RSCP, ASU, BER, CQI, RSSNR, Ec/No.

Limitations arising from the subscription

It is important to recognise when the speed is limited by maximum speed of the subscription rather than the network performance. The network may for example be technically able to deliver a bit rate higher than the purchased subscription, and it's also possible that the speed is throttled to a very low value after reaching the end user's data cap.

Version of the mobile equipment operating system

This item is covered in the fixed environment section above.

Cross traffic

This item is covered in the fixed environment section above.

5.3 Hardware and software information retrieval methods

The aim of this section is to discuss approaches for the measurement client to gather as much reliable information as possible about the presence of factors in the end user environment which might affect measurement results. Note that standalone software/applications are best placed to be able to perform this metadata retrieval functions e.g. apps provided for mobile terminals.

Hardware remote management protocols like TR-064¹⁰ can be used to discover modem router specifications and status and the subscription information the ISP might have provisioned to the modem.

Application programming interfaces (APIs) which are developed and distributed by different software manufacturers (Java¹¹, Microsoft¹², Google¹³, Apple¹⁴...) should be used for programming measurement clients in order to retrieve the metadata. It is worth mentioning that different operating systems may give measurement clients different access for end user environment based factors.

When collecting this information, the applicable privacy rules must also be taken into account. Certain information may constitute personal data as defined in Regulation (EU) 2016/679 and the consent of the end user is required.

5.4 Measurements data filtering

NRAs should retrieve and store all relevant measurement results and available associated environmental information to enable analysis of the effect of end user environmental factors upon the measurement results, to allow a better assessment of results.

The assessment of the measurement results is further discussed in Chapter 6.

6. Measurement results assessment

This chapter provides recommendations for validation of the collected measurement results. This chapter also provides recommendations on how the speed measurement results should be assessed in comparison to the contractual speed values for end users. Guidance is also given on the number and distribution of the individual measurements.

Finally, the topic of data aggregation for market level assessment purposes is discussed and guidance on monitoring the general IAS quality (IAS as a whole and effect of specialised services on IAS) as well as individual applications using IAS is provided.

In assessing cross-border measurements it is important to understand how the server location can impact the measurement results. It is evident that delay increases with the number of network hops, and that connectivity may be affected by crossing multiple network borders. Ultimately the measured quality is dependent on transmission delay, all the links between the measurement client and server and also the measurement server's capacity. While these factors should always be taken into account, this is particularly the case when performing cross-border measurements.

¹⁰ <https://www.broadband-forum.org/standards-and-software/technical-specifications/technical-reports>

¹¹ <https://docs.oracle.com/javase/8/docs/api/>

¹² [https://msdn.microsoft.com/en-us/library/windows/desktop/ff818516\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ff818516(v=vs.85).aspx)

¹³ <https://developer.android.com/reference/android/telephony/TelephonyManager.html>

<https://developer.android.com/training/basics/network-ops/managing.html>

<https://developer.android.com/reference/android/net/ConnectivityManager.html>

¹⁴ <https://developer.apple.com/library/content/referencelibrary/GettingStarted/DevelopiOSAppsSwift/>

6.1 Data validation

The measurement tool generates a data set for each measurement period that is stored in a database. Depending on the kind of measurement approach chosen, the data validation is likely to be complex and extensive. There are two basic measurement approaches:

- a) Measurement campaigns using measurement systems with dedicated clients and servers in a controlled environment or
- b) Crowdsourced measurement campaigns relying on end user initiated measurements using end user equipment.

For the measurement approach based on dedicated clients and servers, simple plausibility checks like time stamps matching the measurement schedule, correct client identification etc. are sufficient, since client and server use dedicated hardware correctly attached to the IAS with software properly installed.

For crowdsourced measurement approaches, more extensive steps should to be taken since the conditions at the client side are not predetermined, i.e. it is unknown whether the client environment fulfils the requirements for an accurate measurement. To some extent this can be cross-checked by the use of metadata (see section 5.2). Also crowdsourced measurements often involve information provided by the end user about the IAS offer being measured, geolocation etc. This information should be validated where possible.

The validation process of end user provided information is a multi-step process, starting with the removal of implausible data. Verifying internet service provider identification, such as rejection of providers which are not relevant for the measurement campaign.

Cross-checking for the correct measurement set-up is done by the use of metadata as described above. Depending on the end user environment requirements, certain metadata should be collected together with each measurement result. Such records could include terminal equipment connection (e.g. Ethernet, Wi-Fi), type of terminal used, status of terminal equipment (e.g. processor load, cross-traffic, parallel active applications), network environment (firewall) or kind of access technology of IAS (e.g. identifying modem type) etc. as described in section 5.2.¹⁵

6.2 Speed assessment for end users

According to the Regulation ISPs must declare the minimum, normally available, maximum and advertised download and upload speed in their fixed network contracts. For mobile network subscriptions ISPs must declare estimated maximum and advertised download and upload speeds.

6.2.1 *Minimum speed*

The minimum speed is the lowest speed that the ISP providing fixed network IAS undertakes to deliver to the end user, according to the contract. The actual speed should not be lower than the minimum speed, except in cases of interruption of the IAS.

¹⁵ When collecting and processing information, the applicable privacy aspects rules must be taken into account

The minimum speed value defined in the contract should be compared individually for results of each measurement calculated as described in section 3.1.

The minimum speed requirements and recommendations apply for fixed networks only.

6.2.2 Maximum speed and estimated maximum speed

According to the BEREC NN guidelines the maximum speed is the speed that an end user could expect to receive at least some of the time (e.g. at least once a day).

The maximum speed value in fixed networks defined in contract should be compared individually for results of each measurement result calculated as described in section 3.1. It is important to compare the maximum speed value against a measurement result and not individual samples within the measurement task or within multiple measurement tasks.

The estimated maximum speed represents an indication of the speed to be expected when using the IAS within the area of coverage of the provider's mobile network. This includes also various use conditions like ideal free-field conditions, in house, in motion etc. Thus verification of the contracted maximum estimated speed cannot be done by single (sporadic) measurements. NRAs may elect to use crowdsourced data and/or drive tests for this purpose.

This recommendation applies for both fixed (maximum speed) and mobile IAS (estimated maximum speed). Note that this recommendation does not specify how often or how many times the measured speed must reach the maximum contractual speed value to confirm that the delivered speed fulfils the contractual promise.

6.2.3 Normally available speed

According to the recitals of the Regulation the normally available speed is understood to be the speed that an end user could expect to receive most of the time when accessing the fixed Internet access service. BEREC considers that the normally available speed has two dimensions: the numerical value of the speed and the availability (as a percentage) of the speed during a specified period, such as peak hours or the whole day.

According to the BEREC NN guidelines the normally available speed should be available during the specified daily period. For example an NRA may set a requirement that the normally available speed should be available at least during off-peak hours and 90% of time over peak hours, or 95% over the whole day.

The normally available speed should be calculated based on measurement results as described in section 3.1.

6.2.4 Advertised speed

According to the BEREC NN guidelines the advertised speed for a mobile IAS offer should reflect the speed that the ISP is realistically able to deliver to end users.

NRAs could set requirements in accordance with Article 5(1) on how speeds defined in the contract relate to advertised speeds, for example that the advertised speed should not exceed the maximum speed defined in the contract.

Whether the advertised speed is realistically reachable for mobile IAS or not should be evaluated on market level rather than individually for each end user.

6.3 Market level aggregation

After the measurement data is processed, the results can be aggregated. At the market level, the measurement results are summarised into aggregated values for different categories such as IAS offers, ISPs, geographical area or similar.

Aggregated results of IAS performance at the market level may be used for regulatory supervision, including monitoring of the general IAS quality.

The market level data can also be used for transparency purposes by publishing statistics, as well as interactive maps showing fixed IAS performance or mobile IAS performance in a geographic area

General considerations regarding how to collect measurement data by the use of crowdsourced measurement approaches, and discussion of advantages and disadvantages of this approach, is provided in BEREC Report on Monitoring QoS of Internet access services in the context of net neutrality, see chapter 4.5.2 [4].

6.3.1 Monitoring the general IAS quality

The performance of an IAS as a whole consists of packet transfer performance (see chapter 3). The KPIs referred to in section 4.2 can also give indications about a poor quality of the whole IAS.

Ideally, monitoring general IAS quality requires the permanent and continuous collection of measurement results. Data which has been generated by crowdsourced measurement tools and aggregated according to 6.3 should be processed and evaluated on a regular basis to enable the monitoring of general IAS quality, ref. Article 5(1). It is also possible to run specific measurement campaigns as required.

The market level aggregated measurement data could be used to monitor that the average available quality, e.g. speed, delay and packet loss of IAS improves over time. In addition it is important to assess whether an ISP e.g. treats individual applications differently (see section 6.4).

6.3.2 Effect of specialised services on IAS

According to the Regulation specialised services (SpS) shall not be provided to the detriment of the availability or general quality of internet access services for end users. Therefore the task for NRAs is to check that specialised services are not provided at the expense of IAS.

According to the BEREC NN guidelines there are multiple approaches as to how NRAs can supervise this, NRAs could for example:

- Request information from ISPs regarding how sufficient capacity is ensured and at which scale the service is offered (e.g. networks, coverage and end users),
- Assess how ISPs have calculated the additional capacity required for their specialised services and how they have ensured that network elements and

connections have sufficient capacity available to provide specialised services in addition to any IAS provided.

- Perform measurements of the IAS

An NRA could assess the aggregated IAS QoS measurement results before and after the introduction of a certain specialised service. If the measured speed values are in general lower after the SpS introduction, this could be seen as an indication that the SpS is provided at the expense of IAS. The NRAs may monitor this e.g. by following the trend how the average speed measurement results for each ISP are evolving. When the introduction of a SpS affects the general quality of IAS, this could be visible also from general IAS performance results.

Another, more direct approach requires the NRA to take into account the network topography supporting the IAS. In this scenario the performance of the IAS is measured while a nearby end-user is using a specialised service and again at a time when no specialised service is being used.

The following is an illustrative example where we consider two consumers on a fixed line IAS:

- Neighbour A: Is an IAS user without specialised services
- Neighbour B: Is a user of both IAS and an IPTV specialised service

It is assumed that both Neighbours are on the same access network segment and/or IP traffic for both neighbours shares a path to the network edge (E.g. shared backhaul, shared DSLAM/CMTS).

In order to detect whether Neighbour B's usage of specialised services is affecting neighbour A's IAS, it is recommended that the throughput of the IAS delivered to Neighbour A is measured before Neighbour B commences an IPTV session. This throughput can be measured again later when Neighbour B is using the IPTV service.

The results can be compared to verify that the use of a specialised service does not impact upon the IAS of other users.

6.4 Individual applications using IAS

Measurements of the performance of individual applications (see Chapter 4) may show whether blocking or any kind of prioritisation or throttling of specific applications is applied to an IAS offer.

Some of these traffic management practices may only be detectable when the network is congested which will require distributed measurements over time in various network segments.

Tools for detecting traffic management practices are likely to provide an indication of the presence of such a practice rather than a clear result. For example, when differences are observed in the calculated weight of a web page (number of bits that are transmitted during the page load) that has been loaded in similar conditions between different ISPs, it could be an indicator to detect blocking of a part of a web site (e.g. ads) or data compression.

When differences are observed for an ISP over time or between different ISPs and while these differences do not match with the overall development of internet access service quality the ISP is offering, it could be an indication that some traffic management practices are applied. For example, the streaming platform can be throttled or prioritised. It could also be useful to have screenshots or records from the measurement tool of the tested videos to assess whether or not there is a difference of quality that could be due to compression for example.

Another way of detecting traffic management practices could be to compare the measurement results related to a specific ISP both with and without a VPN. The key idea is to use a VPN proxy located near the IAS provider's network edge to record and replay the network traffic generated by arbitrary applications, and compare it with the network behaviour when replaying this traffic outside of an encrypted tunnel. If there are significant and recurring discrepancies, it could be a strong indication that there may be impact from traffic management.

Most of the time, a measurement at the application level can only detect the presence of an inadmissible traffic management but not the cause or responsible network segment.

7. Certified monitoring mechanism

According to the Regulation ISPs must describe the minimum, normally available, maximum and advertised download and upload speed in their fixed network contracts. For mobile network subscriptions ISPs must describe estimated maximum and advertised download and upload speeds.

The Regulation 2015/2120 [1] defines that an end user may use a monitoring mechanism certified by the NRA to check that the actual performance meets what has been specified in the contract. This measurement information can be used for triggering the remedies available to the consumer in accordance with national law.

This entails a decision on whether the subscription meets the different speed values defined in the contract and whether there is a significant discrepancy, continuous or regularly recurring. Note that in some Member States the NRA may not be competent to resolve disputes between consumers and undertakings providing electronic communications services, including deciding on whether there is significant discrepancy, and such decisions may be made by a different authority or body.

To be able to issue a declaration that there either is no significant discrepancy between actual and indicated performance or that there is such a discrepancy empowering the user with the right to trigger "the remedies available to the consumer in accordance with national law", a number of conditions should be satisfied from a regulatory point of view for giving legal value to this "evidence". The final ruling over which "evidence" is sufficient for triggering legal consequences however is still subject to court rulings. Therefore, decisions of NRAs should be made transparently; all measurement data should be available for further legal considerations of the respective court.

The Regulation does not require Member States or an NRA to establish or certify a monitoring mechanism. Therefore it is worth noting that a certified monitoring mechanism may be available only in some member states. The Regulation does not define how the certification

should be done, so this is a national matter. If the NRA provides a monitoring mechanism for this purpose it should be considered as a certified monitoring mechanism according to Article 4(4) of the Regulation. As the Regulation talks about a monitoring mechanism certified by the NRA, the question of when to certify a monitoring system and how to certify can be considered to be up to an NRA according to the national legislation and circumstances.

7.1 Guidance on criteria regarding certified monitoring mechanism

This section gives guidance on criteria NRAs could take into account when providing its own certified mechanism or certifying a third party mechanism in accordance with the Regulation and BEREC NN guidelines [3].

- a) The certified measurement mechanism should fulfil the requirements specified in chapter 3 and take the considerations of chapter 5 into account.
- b) The certified monitoring mechanism should be in compliance of the applicable legislation such as privacy rules.
- c) End users should be enabled to make a straightforward comparison between measurement results and the contractual speed values.
- d) The NRA is recommended to give guidance on in which cases a significant and continuous or regularly recurring difference is established by the certified monitoring mechanism. Noncompliance on a single indicator is sufficient to give the user the right to use “the remedies available to the consumer in accordance with national law”.
- e) The NRA is recommended to ensure the integrity of the operation of the certified mechanism in case the mechanism is provided by a third party. It is also recommended to take into account the independence and business model of the entity providing the monitoring mechanism where it is not provided by the NRA itself.

8. References

[1] Regulation (EU) 2015/2120 of the European Parliament and of the Council of 25 November 2015 laying down measures concerning open internet access and amending Directive 2002/22/EC on universal service and users' rights relating to electronic communications networks and services and Regulation (EU) No 531/2012 on roaming on public mobile communications networks within the Union,

<http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32015R2120&from=EN>

[2] BEREC Guidelines for quality of service in the scope of net neutrality, November 2012, http://berec.europa.eu/eng/document_register/subject_matter/berec/regulatory_best_practices/guidelines/1101-berec-guidelines-for-quality-of-service-in-the-scope-of-net-neutrality

[3] BEREC Guidelines on the Implementation by National Regulators of European Net Neutrality Rules, August 2016,

http://berec.europa.eu/eng/document_register/subject_matter/berec/regulatory_best_practices/guidelines/6160-berec-guidelines-on-the-implementation-by-national-regulators-of-european-net-neutrality-rules

[4] BEREC Report on Monitoring QoS of Internet access services in the context of net neutrality, September 2014,

http://berec.europa.eu/eng/document_register/subject_matter/berec/reports/4602-monitoring-quality-of-internet-access-services-in-the-context-of-net-neutrality-berec-report

[5] BEREC Report on Feasibility study of quality monitoring in the context of net neutrality, November 2015.

http://berec.europa.eu/eng/document_register/subject_matter/berec/reports/5576-feasibility-study-of-quality-monitoring-in-the-context-of-net-neutrality