

Notice No. 2015-0832
from the *Autorité de régulation des communications électroniques et des postes*
(Regulatory Authority for Electronic Communications and Postal Services),
of 7 July, 2015,
on the structure of bandwidth usage on internet access networks
in France

The *Autorité de régulation des communications électroniques et des postes* (hereinafter “ARCEP”),

Pursuant to the Postal and Electronic Communications Code (*Code des postes et des communications électroniques*) and in particular article L. 135;

Pursuant to the letter of 9 April 2015, in which the Minister of the Economy, Finance and Industry asked ARCEP to give its opinion on the structure of bandwidth usage on internet access networks in France, and on methods available for its measurement;

Having deliberated on the matter on 7 July, 2015,

Has formed the following opinion:

In the letter of 9 April, 2015, received by ARCEP on 13 April, 2015, the Minister of the Economy, Finance and Industry sought the advice of ARCEP on the “*current structure of bandwidth usage on internet access networks in France, and on the possibility and means of collecting information that would enable detailed knowledge of the nature and source of these traffic streams [...]*”.

The Authority was in particular asked about (i) the relevant technical characteristics for the measurement of bandwidth usage, (ii) the data available about online consumer services established in France that use the most bandwidth resources, (iii) what methods can be used to ensure that these data are collected and improved and, finally, (iv) what strategies may be implemented by online service providers that could complicate or impair the collection of such information.

I Introduction

The development of the internet over the past two decades has turned it into a link between several billion human beings and a factor of social and economic development. As an open global platform, the internet today constitutes a structural common good, whose proper functioning is a crucial issue.

In order to prepare the opinion given in the present notice, the Authority conducted a series of meetings with stakeholders in the sector, which led it to audition eight key internet actors in

France¹. A questionnaire was also sent to several actors seeking, in particular, to obtain technical and quantitative information that could enhance the opinion.

This opinion therefore constitutes a technico-economic contribution by ARCEP to the Government's reflections on the possibility of establishing a reliable and accurate analysis of the traffic flow within internet access networks in France. It does not encompass future possible use that might be made of the results of those measurements.

ARCEP's considerations are based on two essential - and to a certain extent interdependent - characteristics of a traffic measuring system:

- **The depth of analysis** - from simply counting the flow volume to identifying the origin or the nature of the traffic;
- **The measurement point** – the point in the network where the measurement is taken.

In general, the Authority considers that the relevance of a system for measuring bandwidth usage on internet access networks must be assessed on the basis of several criteria: technical and legal feasibility; cost and complexity of implementation; reliability; completeness; and the easily verifiable character of the results obtained.

After several preliminary considerations on the functioning of the internet (section II), the different points within the network where traffic measurement may be carried out and the different depths of analysis that can be envisaged are presented and evaluated, with a comparative analysis of advantages / disadvantages (section III). This is followed by a section on the impact of different delivery modes of internet traffic and the bypass measures that can be put in place by actors who do not want their flows to be measured or identified (section IV).

II Preliminary considerations on the functioning of the internet

This section presents some technical considerations on the functioning of the internet, relevant to the measurement of bandwidth usage on internet access networks.

For more specific information about interconnection arrangements, the major categories of actors and the different internet traffic delivery modes, please refer to the Authority's 2012 report to the Parliament and the Government on Net Neutrality².

II.1 Traffic and growth trends

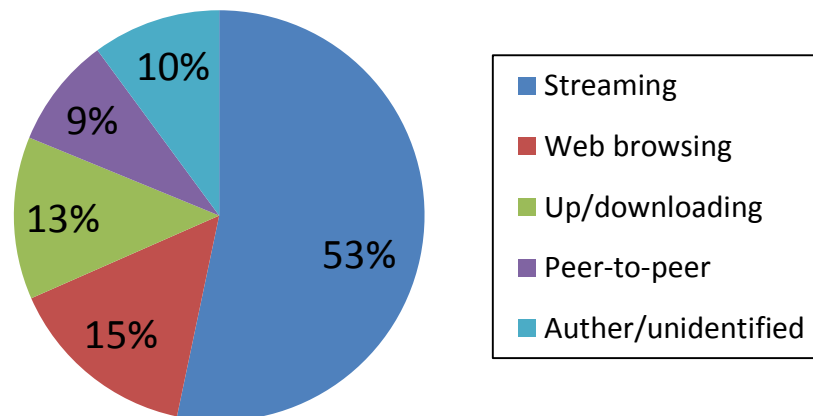
With the development and generalisation of high and very high-speed broadband internet access networks, the number of actors offering innovative uses and services has multiplied. These new uses have been accompanied by an increase in exchanged traffic, growing at a sustained pace, which has, in turn, encouraged the development of more efficient networks. At the global level, it is estimated that over the next few years this traffic will continue to

¹ Two internet service providers (ISPs), two content and application providers (CAPs), an internet exchange point (IXP) operator, an international transit provider, a hosting service provider and a company specialising in the analysis of network performance.

² In particular annex 4, p. 84 and annex 6, p. 93.

grow year on year,³ by nearly a fifth on fixed networks and by half on mobile networks. The trends in France should broadly follow those for the global level.

The nature of the exchanges that take place on the internet has also changed profoundly. In proportion, web browsing has become largely a minority activity, supplanted in particular by the dissemination of audiovisual streams, which consume a large amount of bandwidth. Video streams are experiencing such a boom that, according to the operators, they now represent more than half of the traffic carried and have become the main driver of the growth in traffic.



Usage distribution across all carried traffic⁴

This evolution is accompanied by a profound change in paradigm. Although the internet was originally developed in a decentralised way, on the basis of direct exchange between users, the last decade has seen the development of usage, particularly in the audiovisual field, based on a client-server model.

Furthermore, the traffic has gradually become concentrated around a few major sources, resulting from the merger of the main content and application providers (CAPs) that have emerged at the global level. For example, the information provided to ARCEP during preparation of this opinion reveals that more than half of the traffic carried on internet networks in France is delivered by only 5 autonomous systems⁵ (corresponding to service suppliers, hosting providers or technical intermediaries), of the 60,000 that make up the internet.

These two phenomena of constant increase and concentration of traffic have caused a major change in the traffic profile, which now moves predominantly downwards, from the main global CAPs towards internet users.

³ Source: CISCO - *Visual Networking Index Forecast*, February 2015.

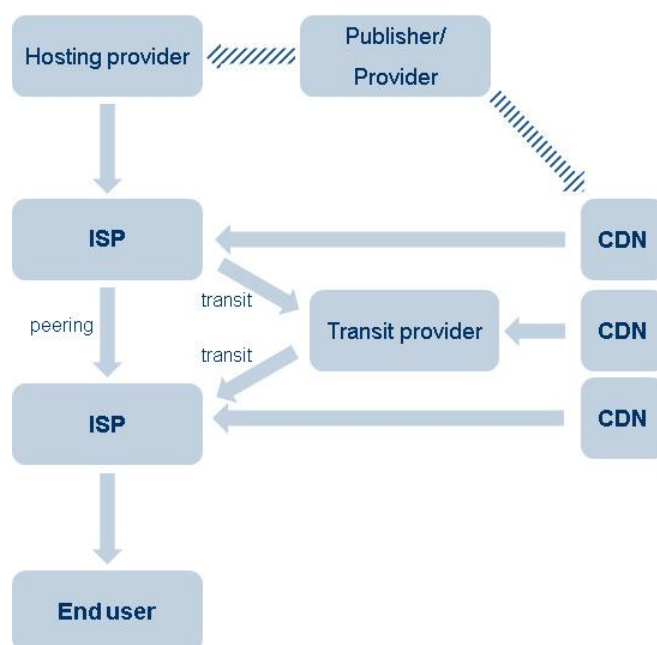
⁴ Source: replies to the questionnaire sent out by ARCEP.

⁵ An autonomous system (AS) is the smallest level of internet unit, and designates a network, or a set of IP addresses, under the control of a single entity.

II.2 Traffic delivery methods

To keep up with the changes described above, internet traffic delivery methods have undergone a transformation. Notably, various types of technical intermediaries have emerged, dedicated to improving and optimising the dissemination of audiovisual content, as well as technical solutions aimed at bringing content closer to the end user by means of cache or CDN servers⁶ (which host a copy of the content).

These delivery modes can be combined, as shown on the following diagram.



Internet traffic delivery methods (illustrative example)

II.3 Difficulties related to the identification of “online service providers”

The internet is characterised by a strong decorrelation between, on the one hand, the content, services and applications provided and, on the other hand, the networks that ensure their dissemination, in a predominantly undifferentiated manner.

In this regard, it should be emphasised that the notion of “online service provider” mentioned in the letter of referral is not, as such, a legally defined term.

Furthermore, while, as the Council of State notes in its 2014 annual survey, *Le numérique et les droits fondamentaux* (Digital technology and fundamental rights), the Trust in Digital Economy Act formulates a *summa divisio* between website publishers and technical intermediaries such as hosting providers, which is based on the idea that the latter have a purely passive role without knowledge of the information they store, it is more and more difficult to make this distinction in practice.

⁶ Content delivery networks.

By way of illustration, for a same platform of online videos, it can be complex to identify the content publisher, which could be, depending on the case, the platform operator or a third party.

Taking into account the uncertainties mentioned above, and independently of any technical considerations, there is therefore no single, unequivocal, infallible or exhaustive method for associating an IP packet circulating on the internet with an online service provider or, respectively, to identify the set of IP packets associated with a given online service provider.

On the technical level, a thorough analysis of traffic flows and their content (through *Deep packet inspection* or DPI, see below) could permit an in-depth knowledge of the nature of a service. A case-by-case analysis may then make it possible to identify the provider or providers of relevant services, depending on the definition adopted for this concept. However, the increasing use of data encryption - which aims to ensure the integrity and confidentiality of information exchanges - can make this analysis inoperable in practice.

According to the information gathered from different actors by ARCEP during the preparation of this opinion, in mid-2015, the share of encrypted traffic on the internet as seen by ISPs⁷ in France runs at nearly 50 per cent, as against approximately 5 per cent in 2012⁸.

III Measurement and identification of the characteristics of internet traffic

In this section, after introducing several aspects of metrology (III.1), the different levels of analysis according to which measurements can be performed are presented and compared (III.2), followed by the various points in the network where a measuring system can be implemented (III.3).

The advantages and drawbacks of each method are evaluated in light of the criteria set forth in section I.

III.1 Some aspects of metrology

III.1.1 Relevant volumes

In order to scale their networks (links and active equipment), operators most often base themselves on a bitrate indicator (usually expressed in Gbit/s) of traffic streams routed over one or several given links. This indicator reflects the amount of data passing through a link or through the entire network per unit of time. It can therefore take account of possible peaks of traffic due to an increase in usage. In practice, this is how an increase in capacity is planned for and then implemented by the operator, when the observed bitrate reaches a certain proportion of the capacity on a part of its network. The bitrate indicator therefore reflects the investments made by the access providers.

In turn, the volume indicator, in terabytes⁹ (TB), translates the total amount of data having flowed through the network during a given period. Although providing some insight on end-user consumption, it appears less relevant for the quantification of bandwidth usage. In fact,

⁷ Internet service provider.

⁸ Source: Response to the questionnaire sent by ARCEP.

⁹ 1 Terabyte = $8 \cdot 2^{40}$ bits $\approx 8 \cdot 10^{12}$ bits.

without reference to the dimension of time, it does not enable an operator to correctly size their network and is therefore not representative of their recurring costs or investments undertaken by them. It is in general not tracked by ISPs on fixed networks.

III.1.2 Sampling frequency and reference value

The traffic sampling frequency corresponds to the time interval between two measurements. The higher the sampling rate, the more correctly the measures can represent the rate of traffic carried. However, considerable storage and processing capacity would be necessary to manage the volume of data generated. A low - and less expensive - sampling rate risks missing the sudden peaks or drops characteristic of the internet traffic profile.

In order for it to be representative of bandwidth use on a network, the reference value used in the framework of a measuring system should include high demand periods of the day.¹⁰ Including those periods during which there is a heavy call on bandwidth resources as sampling reference points allows an ISP to ensure adequate routing capacity, including during times of increased traffic.

According to the information gathered by ARCEP, the most common method used by operators, especially for the purposes of sizing or for billing certain traffic delivery services, is to measure the volumes carried at five-minute intervals, and to deduce the average bitrate based on these intervals. The daily or monthly reference value used is then the 95th percentile of these flows. This value represents the bandwidth usage of a network during a period of very high load, whilst at the same time eliminating the highest peak measurements, which correspond to rare phenomena and are of short duration.




Illustration of traffic (flow) trends during a typical day

For more information on these aspects of metrology, the work of the International Telecommunications Union (ITU) and the Internet Engineering Task Force (IETF) on these questions is particularly recommended.

III.1.3 Traffic asymmetry

Two types of traffic streams can be distinguished within an ISP: outgoing flows, issued by its customers, and incoming flows, traffic that reach them. The development, in the course of these last few years, of services based on the client-server model is today reflected by a

¹⁰ Also called "peak" hours, as opposed to "off-peak", when there is lesser demand for network resource usage.

marked asymmetry between these incoming and outgoing flows. This asymmetry increases as richer content, with a greater bandwidth consumption (for example, high definition audiovisual content), are made available to internet users. The ratio between the outbound and inbound traffic at the interconnection points between ISPs and major content providers can be one to twenty, or even greater¹¹.

Since the transmission links are bi-directional, the reference value for network sizing (or transit agreements, for example) is the highest value between the inbound and outbound flows. When estimating the bandwidth usage of an internet access network in France, it seems possible that outgoing traffic flow may be neglected, relative to incoming traffic flows.

III.2 Characterisation of internet flows: the different levels of analysis

Measurement of bandwidth usage on internet access networks can be carried out at three main levels of analysis. These are, in ascending order of granularity and complexity:

- Recording of data volumes passing through a link by using meters embedded in the network equipment, with data collection via the SNMP protocol. It is possible to deduce the traffic bitrate through regular requests to these meters at frequent intervals;
- Identification of the source and destination of traffic, thanks to tools such as Netflow/IPFix, which allow the addresses indicated in the IP packet headers to be collected, therefore make it possible to deduce the originating AS.
- Analysis of the content of transported packets, thanks to techniques and probes such as Deep Packet Inspection (DPI), which make it possible to retrieve richer information on the nature of the services transported.

¹¹ Estimate produced from the interconnection data collected periodically by ARCEP as part of ARCEP decision no. 2012-0366 on the introduction of information gathering on the technical and pricing terms governing interconnection and routing of data, as amended by decision no. 2014-0433-RDPI.

Level of analysis	Tools used	Advantages	Drawbacks
Measurement of traffic flow	<ul style="list-style-type: none"> - Software solutions that, through the SNMP protocol, gather the information that is continuously collected by the different network devices (switches, routers, etc.) 	<ul style="list-style-type: none"> - Reliable data. - Certifiable and exhaustive information (at network device interface level). - Standardised and commonly used SNMP protocol (RFC 3411 and later). - Inexpensive tools, easy to implement. 	<ul style="list-style-type: none"> - No identification of form of flow (origin, destination, nature of service, etc.). - Identification of the interconnected actor¹².
Identification of the origin and destination (AS) of the traffic	<ul style="list-style-type: none"> - Softwares based on the Netflow / IPFix protocols, enabling retrieval and aggregation of information concerning the flows running through a given router. 	<ul style="list-style-type: none"> - Enables identification of the originating IP address of the flow and deduction of the originating AS (idem for the destination). - Enables identification of the ports used and to deduce from this the use of some wide-spread applications (web browsing, e-mail, etc.). - Standardised IPFix protocol (RFC 7011). - Compatible with a wide variety of routers and systems of administration. - Widely used and relatively easy to implement. 	<ul style="list-style-type: none"> - Statistical results based on flow sampling (current usage consists most often of a sampling rate of 1:1000). - Significant costs and risk of equipment performance degradation in case of over-frequent sampling. - Identification of the originating AS of a flow does not systematically allow the publisher of content to be identified.
Identification of the nature and publisher of the flows	<ul style="list-style-type: none"> - Service cards installable in routers. - Software solutions deployable on network equipment. - Physical probes installable at output ports of network equipment. 	<ul style="list-style-type: none"> - Enables identification of the originating IP address of the flow and deduction of the originating AS. - Enables identification of the nature of the flow and therefore the type of associated use. 	<ul style="list-style-type: none"> - Statistical results based on flow sampling - Risk of equipment performance degradation in case of over-frequent sampling. - Very costly solutions to implement at the level of individual operator networks: important investment and operating costs in storage and processing capacity. - Intrusive nature of the analysis carried out,


¹² In the case of a public peering within an internet exchange point, the data obtained would not enable identification of who, among the members of the exchange point, was originator or recipient of an IP packet.

			<p>which may raise questions in relation to the protection of personal data.</p> <p>- Very degraded results in the case of data encryption.</p>
--	--	--	---

III.3 The different points of traffic measurement

The flows carried by an ISP pass through different levels of its network between the inbound transit point and the end user. Traditionally, four network segments are identified:

- The local access network (local loop), i.e. the section located between the subscriber and the first active network access equipment;¹³
- The transport or backhaul network constitutes the first level of end user traffic aggregation. Its role is to carry traffic between the access equipment and the first routers (BRAS/BNG¹⁴), which manage authentication and access to the service and provide the interface with the core network;
- The core network, the level at which the traffic is carried at national or international level and sometimes to the edge of the network. It is composed of a small number of very high capacity routers, forming a very resilient grid. The core network is the so-called “backbone” of an ISP’s network;
- The point of interconnection, which is made up of internet routers and clusters marking the boundary between the network of the ISP and the rest of the networks that make up the internet.



*Schematic Architecture of an internet access network
(example of a fixed network with copper local loop)*

The traffic can be measured at the interface between each of these layers, which is characterised by a category of active equipment that can accommodate a device to measure the traffic. At each of these points it is, apart from the exceptions listed below, possible to collect measurements at the different levels of analysis and granularity discussed in III.2.

¹³ For example: DSLAM (*Digital Subscriber Line Access Multiplexer*) for copper access networks; CMTS (*Cable modem termination system*) for wired access networks; OLT (*optical line terminal*) for optical fibre access networks; NodeB for mobile 3G access networks and eNodeB for mobile 4G access networks.

¹⁴ *Broadband remote access server / Broadband network gateway.*

	Advantages	Limitations
First active equipment (access) level measurement	<ul style="list-style-type: none"> - The measurements carried out at access level are exhaustive, since all the individual flows are necessarily routed via this segment of the network. 	<ul style="list-style-type: none"> - The number of measurement points is very high (in the order of tens of thousands) which leads to increased implementation costs, particularly in terms of the number of licenses for the measuring tools. - Access equipment is generally dedicated to transmission and multiplexing, it is not adapted to a granular analysis of traffic. The installation of specific equipment might prove to be necessary.
Edge router (aggregation) level measurement	<ul style="list-style-type: none"> - The number of measurement points is reasonable at individual operator scale (in the order of a few hundred). - The measurements are carried out at service access platform level (BRAS, BNG), which aggregate and control the traffic from the end users: all subscriber traffic is therefore routed via these. - It is at the edge router (aggregation) level that we observed the strongest correlation between volume of routed traffic and network costs. 	<ul style="list-style-type: none"> - Certain flows routed in aggregate are only divided into individual streams at the access level and therefore are measured only in a partial manner at the backhaul (aggregation) level.
Interconnection point level measurement	<ul style="list-style-type: none"> - The number of measurement points is reduced (fewer than a hundred). - The monitoring of traffic routed by internet routers is a widespread practice among operators. - Each physical link identifies an interconnected actor (transit operator, internet exchange point, actor interconnected by means of private peering). 	<ul style="list-style-type: none"> - Certain types of flows that are routed in aggregate at the interconnection point and are only divided into individual streams at a lower level in the network, are measured only in a partial manner. - Some flows do not pass through the interconnection point and therefore are not measured.

III.4 Partial conclusion

First of all, the Authority notes that DPI techniques, which provide fine-grained analysis of the nature of flows, must be discarded, without it being necessary to examine more precisely their technical or legal feasibility. This is because, in an environment where data encryption solutions are being more and more widely used (as described in the section II.3), DPI techniques do not enable a sufficiently complete view of the traffic.

A first method based on the Netflow/IPFix protocol and implemented at service access platform level, located on the border of the backhaul (aggregation) network, could be envisaged. This method would allow coverage of more traffic than at the interconnection point, and would enable a first level of identification to be carried out by analysing the origin

and the destination of the flows. However, it is still based on a statistical sampling of the traffic, which does not provide easy verification of results and presents network performance degradation risks in case of widespread and intensive use.

A second method could be to read the meters continuously maintained by the network equipment (through the SNMP protocol). Applied at the level of the interconnection links, this method, while it does not make it possible to identify the origin, destination and the nature of the flows, does make it possible to determine the actor who passed the traffic to the ISP, thanks to the prior knowledge that ISPs have of interconnected partners. It also presents several advantages in terms of reliability of results and above all in terms of feasibility for the ISP. The Authority therefore considers that this method presents the best compromise in terms of the criteria of the analysis presented in section I.

It should be noted, however, that this method does not allow traffic corresponding to content hosted within the ISP network to be measured. To take account of this traffic, measurements at the interconnection level should be supplemented by measurements carried out, using the same tools, at the level of the ISP network internal servers. Further work would be necessary, however, to verify the feasibility of these measurements and to specify the servers to be taken into account.

Finally, it is important to note that the two methods presented above are not complementary. Since each of these devices operates at a different measurement point in the network, their respective results would be difficult to reconcile and might lead to multiple counting of the same flows.

IV Drawbacks and possibilities of bypassing a measuring device

Several conveyance modes exist which should be taken into account when establishing a system for the measurement and characterisation of internet traffic. Some of these modes - already widely integrated into the strategies of operators and CAPs for their traffic routing - can lead to partial measurement of the relevant flows, depending on the measuring points selected.

Additionally, certain practices on the internet that aim to optimise bandwidth usage or to secure data exchanges may significantly reduce the ability of a measuring system to identify the origin, destination or nature of the flows. A stakeholder who did not want its traffic to be measured or identified could therefore use these as a bypass method.

In the context of this opinion, the Authority would like to point out several practices that can restrict the useful scope of a measuring device:


- Content hosting within the ISP network;
- The development of peer-to-peer exchanges (i.e. the hosting of content by end users) ;
- The localisation of interconnections outside the national territory;
- Multicast distribution

This section will therefore consider the impact that various routing modes and possible bypass techniques can have on the ability of a tool to obtain comprehensive and reliable measurements and identify the origin and nature of flows.

IV.1 Content hosting within the ISP network

Internal servers such as cache servers or on-net CDNs are platforms located within an ISP network, used to store local copies of certain content available on the internet. They aim to bring content closer to the users, in order to optimise cost and performance: their use thus reduces the use of interconnections. With this mode of transport, higher-audience content is routed only once through the interconnections to feed¹⁵ these servers, which will subsequently serve users as many times as the content is requested.

According to the information gathered by ARCEP during the preparation of this opinion, it may be considered that a standard ratio between incoming and outgoing traffic for a cache server or a CDN is approximately 1:5. The use of cache servers or internal CDNs therefore enables a reduction of the use of interconnections of almost 80 per cent for content stored on them.



Internal cache server operation


A measuring device operating at the level of the ISP's interconnection points would account for the flows associated with each content only once, without being able to take into account the multitude of individual flows generated downstream in the network. The same goes for the content requested by the subscribers of an ISP who are hosted by the ISP itself: their traffic will not transit via the interconnection.

In order to ensure completeness of measurement in the case of content hosted within the ISP's own network, it is therefore necessary to place the measurement point downstream of the point of dissemination.

IV.2 Content hosted by end users: peer-to-peer

Peer-to-peer or P2P is a technique that allows data to be exchanged between Internet users, as opposed to the client-server model, which is presently the more prevalent architecture on the internet. Specifically, each user becomes a content supplier, making content available to the other users of the service. This mode of transmission is used in the context of file-sharing applications or even in distributed computing and provides a solution to the availability and load issues inherent in centralised models.

¹⁵ Caching



Peer-to-peer traffic routing (between subscribers of a same ISP)

Over the last few years, service publisher's interest in peer-to-peer type applications has been revived, especially for file sharing applications. In fact, in the context of bandwidth usage optimisation and routing cost reduction, several service providers are increasingly opting for this architecture.

Being exchanged between end users, the traffic associated with exchanges between the subscribers of a same ISP remains internal to the network. A measuring device operating at the interconnection level will therefore have no visibility of this traffic.

In addition, in the case of peer-to-peer exchanges, a tool that aims to identify the flows would designate the ISP customer who hosts the shared content as being the origin of the associated traffic.

IV.3 The localisation of interconnections outside the national territory

Where a measuring device is implemented at the level of interconnection points, it is important to ensure that the legal basis for this tool does permit the collection of information from interconnections originating outside the national territory.

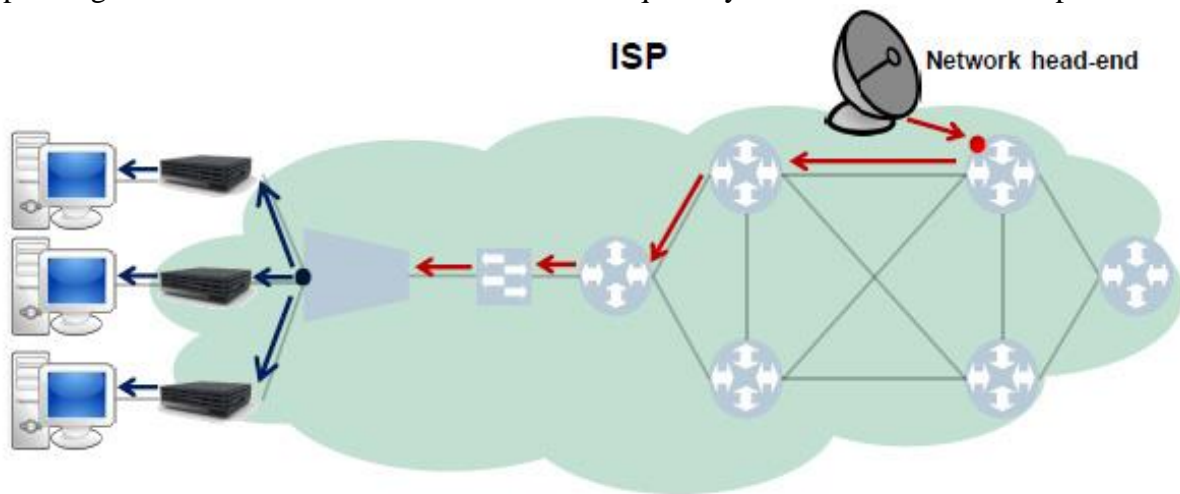
By default, some stakeholders might be incited to move their interconnections or to host their content abroad, which would have an adverse impact on investment and infrastructure deployment in France. This sort of network delocalisation could also degrade the quality of service experienced by users located in France, although this point would need to be assessed in the light of the nature of the services.

IV.4 Multicast distribution

Multicast is used for the distribution of part of (linear) television services on IP, called "*IPTV*". It allows a provider, after retrieving and transforming the signals at the network head-end, to route them in aggregated form for as far as possible, usually up to the last access equipment before the subscriber's line. Division into individual flows is carried out only at this level of the network, according to subscriber demands.

This mode of routing aims to optimise bandwidth usage on the networks by avoiding unnecessary division of identical individual flows (simultaneous viewing and broadcasting), as opposed to, for example, non-linear video services (individual viewing and broadcasting), which are generally disseminated in *unicast* (one flow end-to-end per subscriber).

In the context of multicast, the traffic is transmitted from the operator network head-end and the interconnection points at the network border are not used. A measurement system operating at this level will therefore not be able to quantify the related final consumption.



Multicast distribution

Similarly, a device operating at the backhaul level will measure only the aggregated feed without being able to take into account the traffic division which occurs downstream in the network¹⁶.

However, from the data collected through its questionnaire, the Authority considers that aggregated multicast streams routed from the network head-end up to the access point, passing through the backhaul level, would only represent a bitrate of 1 to 2Gbit/s, as opposed to several Tbit/s, when all the individual flows consumed by subscribers are aggregated. This bitrate would represent more than half of the bandwidth available at end-user access level.

¹⁶ Cable networks being an exception, where the audiovisual streams are introduced downstream of the backhaul and are not therefore measured at this level.

V Conclusion

Measuring internet traffic is a task made complex by the diversity of modes available for its delivery. IP packets may indeed take a variety of different, more or less direct paths to reach the end user, as part of a single exchange. These paths travel over multiple networks and can rely on cache servers – where a version of the requested content is stored – that are located more or less close to the end user.

Moreover, there is no single, unequivocal, infallible or exhaustive method for associating a traffic stream with a service. The nature of the traffic is indeed multifarious – e.g. live and on-demand video, telephone calls, files, websites, etc. – and one characteristic feature of the internet is the strong decorrelation between content, services and applications on the one hand and the networks used to carry them in a way that is largely undifferentiated, on the other. Identifying the type of service being carried therefore requires a very detailed analysis of the streams, which could not only raise concerns in the area of data privacy, but may be rendered inoperable by the encryption of the data being transmitted.

These drawbacks are intrinsic to how the internet operates, but also constitute its strength since they ensure both its resiliency and its openness.

It was with all this in mind that ARCEP assessed the different options for measuring internet bandwidth usage. These options were evaluated based on the following criteria: technical feasibility, cost and complexity of implementation and accuracy, completeness and verifiability of the collected data, while complying with the principles of internet neutrality and confidentiality of correspondence.

In terms of depth of analysis, the Authority concludes that measuring the volume of data traffic on a given link thanks to SNMP would make it possible to obtain accurate and verifiable results, and be relatively easy to implement, compared to other forms of analysis such as Netflow/IPFix or Deep Packet Inspection (DPI).

As to the point in the network where the measurements should be taken, ARCEP concludes that applying this method (i.e. measuring volume using SNMP) at interconnection points would prove a reasonable undertaking for internet service providers (ISP) from a technical and economic standpoint.

This type of measurement would have two drawbacks, however:

- Firstly, because it involves merely counting the number of packets transiting over the network, the stream is identified as coming from the interconnection partner. Generally speaking, this would be a technical intermediary (transit operator, Content Delivery Network, Internet Exchange Point (IXP) manager, etc.) and not a service provider;
- Secondly, because measurements are performed at interconnection points, they do not take into account (in any event not directly) the volumes of traffic being relayed under special conditions such as multicasting, peer-to-peer exchanges between two subscribers with the same ISP and transmissions from an ISP's own hosting centre or a server on its own network.

As a result, there is also a risk that implementing such a measurement system would lead certain players to adopt bypass techniques. Except for peer-to-peer traffic, which cannot be

associated with a specific service or service provider, one path to explore in order to limit this sort of bypassing would be to supplement measurements made at the interconnection level by measurements carried out by the ISP on specific links in the backhaul network, particularly as regards internal cache servers.

Done in Paris, 7 July, 2015

The Chairman

Sébastien SORIANO