27th July 2012

FP7 ETICS Researchers' Response to the BEREC Consultations on the Net Neutrality.

Documents:

- "An assessment of IP-interconnection in the context of Net Neutrality - Draft report for public consultation" - BoR(12)33 [1];
- "Differentiation practices and related competition issues in the scope of Net Neutrality - Draft report for public consultation" - BoR(12)31 [2];
- "Guidelines for Quality of Service in the scope of Net Neutrality - Draft for public consultation", - BoR(12)32 [3].

Executive summary

The debate on the Network Neutrality rose as a public debate from [1], evolving into a Net Neutrality debate. Such condensed wording however leaves room for interpretation, as "Net" could both refer to Network or to Internet. In the meaning of "Internet neutrality", all means to find information (including search engines), store it (including caching) and transport it should be considered.

Although BEREC reckons such a notion of "Internet neutrality", the documents essentially focus on the telco operators (also called Internet Service Providers, ISPs) and argue against traffic management practices. Such a restriction intrinsically induces to look at the aforementioned problem through a given prism, which necessarily influences the conclusions of the analysis. In our response to the BEREC consultation, we advocate that **for a deeper analysis of the issues and for the defence of an open Internet, a wider scope must be considered, and should be debated under the Internet Neutrality idea**.

As telco experts involved in the European Research funded ETICS project, we perceived mainly strong statements about the negative side effect of traffic management and QoS interconnection, while the current markets data and their evolution over past years, as well as the scientific literature on the subject are inconclusive.

Therefore, in our response, we aim to **counterpoise some of the arguments** presented in BEREC's documents analyzing the current situation both from a market and scientific/research view point. We thus argue that:

- Taking into account the new equilibrium of the Internet ecosystem and in the context of a wider debate on the Internet Neutrality, we recommend to not oppose and have a similar treatment between integrated CAP-driven, CDN-driven or ISP-driven possible innovations to improve the future Internet sustainability for all actors, as well as social welfare for Internet' users.
- It is demonstrated from the OTTs' innovations that quality improvement is required for certain applications due to the unpredictable behaviour of the network. In the documents, general conclusions are drawn from specific current services for which stop-gap solutions exist, to argue that QoS differentiation in the network is not required. This is not appropriate and may be detrimental for future advanced services with more stringent constraints.
- More fundamentally, most authors acknowledge that today, de-ossifying the Internet economics is a prerequisite to any technical or service innovation thus meaning that the traditional reachability service has attained its limits in this field.
- As discussed by researcher and experts in the current literature on the subject, no consensual conclusion has emerged analyzing the positive and negative impacts of QoS management. It can therefore be observed that network QoS could foster

market opportunities for many actors of the value chain and thus re-balance cost sharing.

As an illustrative example, possible solutions from the ETICS project are mentioned as a new interconnection framework, allowing for more sustainable relations between actors in the Internet Ecosystem and more choices for Internet users in terms of possible experiences.

0. Introduction: who are we?

This document provides some comments on the BEREC public consultations, from experts in the ICT and networking areas, involved for several years in research and development activities on the "future of Internet".

In particular, as experts involved in the FP7 ETICS project[1] (Economics and Technologies for Inter-Carrier Services), we found relevant to provide some comments and facts on the topics of Quality of Service assurance and IP interconnection.

Contributors:

N. Le Sauze, H. Pouyllau, Z. Ben Houidi, (Alcatel-Lucent Bell Labs France)
M. Roettgermann, I. Khortals (Deustche Telekom AG)
A. Miron (Israel Institute of Technology - Technion)
J. Meuric (Orange Labs)
H. Lonsethagen, Finn Tore Johansen (Telenor ASA)
M. Dramitinos (Athens University of Economics and Business – Research Center)
D. R. Lopez (Telefónica Investigación y Desarrolló)

N.B.: Arguments developed in this contribution represent individuals' opinion, as a group of experts, and do not necessarily represent their respective company or institution opinion.

1. Introduction: general comments - Internet Neutrality Vs Net(work) Neutrality

The debate on the Network Neutrality rose as a public debate from [4], evolving into a Net Neutrality debate. Such condensed wording however leaves room for interpretation, as "Net" could both refer to Network or to Internet. The general definition, as cited in BEREC's response to the European Commission's consultation on the open Internet and Net Neutrality, is *"A literal interpretation of Network Neutrality, for working purposes, is the principle that all electronic communication passing through a network is treated equally. That all communication is treated equally means that it is treated independent of (i) content, (ii) application, (iii) service, (iv) device, (v) sender address, and (vi) receiver address. Sender and receiver address implies that the treatment is independent of end-user and content/application/service provider."*[2], provided that the referenced "network" is the public Internet.

We believe that the definition is correct when talking about Network Neutrality, but that it is too limited. Currently, most of the current work, articles, consultations and public debates solely concentrate on the Neutrality of the Network operations. Such a restriction intrinsically induces to look at the aforementioned problem through a given prism, which necessarily influences the conclusions of the analysis. On the contrary, all means, techniques, actors, etc. involved in transporting any kind of information between two hosts in the Internet should be analyzed and evaluated equally. Therefore, all means to find information (including search engines), store it (including caching) and transport it should be considered.

We therefore advocate that for a deeper analysis of the issues and the defence of an open Internet, a wider scope must be considered, and should be debated under the Internet Neutrality idea.

As such, issues for instance like (1) the search-neutrality, i.e. a principle by which search engines should be impartial when presenting their search results, and (2) the content

[1] The ETICS-project is funded by the European Commission through the 7th ICT-Framework Program. Grant agreement: FP7-248567, Number: INFSO-ICT-248567. www.ict-etics.eu
[2] p. 4

providers' vertical integration (playing both roles of content production and transport) should be as well part of the debate.

BEREC is in some way recognizing this wide scope, in [3][3]: *"In this respect, it is important to take into account the entire chain of stakeholders when monitoring the evolution of markets in general".* It is then regrettable that in [2], arguing that this segment is the only regulated part in the Internet, this view is reduced to the network segment. Citing [2][4], *"Upholding the principle of neutrality concerns all of the players involved in the "Internet chain", whether the parties operating electronic communications networks routing internet traffic, or the many and various providers of services in the information society. As such, some of the questions raised in the debate around internet neutrality fall outside the realm of the rules and regulations that apply only to electronic communication networks. These networks nevertheless occupy a central place on the "internet chain" and among the players that populate it."* On the contrary, and as argued further in this answer and even recognized by BEREC in [2], the definition of network operations and the categorization of actors is becoming more and more vague, which should prevent the analysis of partial segments or of a subset of actors.

Instead, the documents essentially focus on the telco operators (also called Internet Service Providers, ISPs) and argue against traffic management practices. Even when the document reckons the benefits of traffic management policies with respect to the social welfare of end-users, these benefits are almost always devalued by side remarks. We cite as an example, in [1][5], *"If a guaranteed (end-to-end) network performance over IP-based networks is desired […] Operators are free to develop this as competitive markets are often built on quality differentiation, which can generally be considered to be welfare-enhancing. Nevertheless QoS traffic classes introduce a potential for anticompetitive behaviour. This relates to the fact that there might only be a willingness to pay for a premium traffic class in case the best effort class quality is "bad enough".* We indeed agree that bad best-effort services may lead to anticompetitive behaviours, but believe that this is the area where the regulator should be involved, rather than theoretically opposing efficiency-enhancing traffic management practices.

On the contrary, others actors, CAPs and CDNs in particular, are presented in a more positive way, ignoring or minimizing their role of data transport and lowering the possible deviations associated to their positions in the overall Internet ecosystem. Citing [2], *"Dynamic impacts consider the incentives to invest and innovate in the different parts of the value chain. In particular, it considers the impact of differentiation practices on the incentives to invest and innovate of the various parties."* [6] Such statements leads to the following narrow-scope conclusions, as cited further down: *"As stated above, due to network effects of Internet, any restriction could create entry barriers either for end-users or, in particular, to CAPs, reducing this virtuous circle and affecting future consumer welfare."* [7] or *"demand for capacity can be accommodated thanks to continuous and steady technical progress in electronic equipment as well as new innovative techniques like content distribution networks (CDNs) and peer-to-peer communication which would be important contributions to enhanced network architecture for high capacity best effort communication"* [8].

Moreover, CDNs and Best Effort network services do have their limitations as admitted also in [5] *"For the future, the growth of multimedia traffic, including delay-intolerant applications such as voice-over-IP (VoIP), will imply a growing need for differentiated quality of service (QoS) to accommodate the requirements of different types of traffic. The lack of QoS support in the legacy "best efforts" Internet has quite possibly hindered the emergence of some*

---

[3] p. 19
[4] p. 7
[5] p. 7
[6] p. 7
[7] p. 8
[8] p. 18

applications that demand enhanced services, and the ability to cache content enabled innovation in content distribution and facilitated the rise of Content Distribution Networks. We can speculate that there may be new sorts of applications that cannot be supported using CDNs or built by exploiting application-level interconnection with other kinds of networks such as the PSTN. In order for such applications to emerge, the community of ASes will need to coordinate and provision better end-to-end services. Tier 1 ASes have begun to implement interprovider QoS over peering links to support VoIP. Interconnections are also emerging for interprovider VPNs, developments that may signal the need for more efficient coordination and contracting mechanisms by ASes, perhaps as packet-transport continues to become a commodity market." It is clear that analysing the impact of Net Neutrality issues on the health of the Internet ecosystem, the evolution of new services and the richness of customer choices is still at a preliminary stage and it is crucial to include multiple non-trivial aspects to any kind of analysis attempt.

The reduction of the scope of the analysis to the Network Neutrality and to the tools used by network operators to manage traffic is further narrowed while introducing the notion of "best-effort neutrality". While in [2][9], it is mentioned that *"This description of Net Neutrality is very close to the situation in a world of widespread pure best efforts, even if best efforts and net neutrality are not exact synonyms."* in all documents the two terms are often used interchangeably, as for example in the Question 20 of [1][10] *"Q20: Do you subscribe to the view that CDNs lead to improvement of QoS without violating the best effort principle?".*

Finally, the document [1] contains strong statements about the negative side effect of QoS and QoS interconnection, while the current markets data and their evolution over past years, as well as the scientific literature on the subject are inconclusive.

For these reasons, the positioning of the documents appears to us as unbalanced and not forward-looking. The goal of this answer is therefore to counterpoise the arguments presented in BEREC's documents analyzing the current situation both from a market and scientific/research view point. This analysis is structured around three axes:

- The actors and their respective market powers in the Internet ecosystem
- The various definitions of QoS and their corresponding needs
- The positive and negative impacts of QoS management

Then, as an illustrative example, possible solutions from the ETICS project are explored as a new interconnection framework, allowing for more sustainable relations between actors in the Internet Ecosystem and more choices for Internet users in terms of possible experiences.

Based on these arguments, we consider that it is not possible to answer the questions included in [1]. We even argue that due to the complexity of the problem, no actor is really in position to fully answer these questions. However, we, as a research group that includes academic and multi actor type commercial contributors, are in a good position to consider different point of views coming from different commercial and academic organizations, and provide balanced and well thought conclusions to this "Net" neutrality debate.

2. Actors and the respective market powers in the Internet ecosystem

2.1 The Internet Interconnection foundations

Up to now, Peering and IP-transit agreements are the dominant charging mechanisms in the best effort public internet, focusing on one traded good, the **reachability** of other hosts. In peering, two partners allow each other to *reach* their respective customers. Peering is the voluntary and free exchange of Internet traffic between two partners for termination

---

[9] p. 5
[10] p. 40

purposes, based on mutual benefit [6][11]. The peering itself is typically settlement free for both parties, but there are of course costs involved for both peering partners. Both parties have to pay monthly fees for the local loop, the rent for collocation space for routers and servers at the Internet Exchange and the acquisition, installation and maintenance of necessary equipment [7][12]. In addition, each peer incurs cost in carrying the traffic of its peer's customers.

As stated by BEREC there are several requirements set out in the peering policies of an ISP, addressing the extent settlement-free peering arrangements come into account [5][13]. One of the most important conditions is the symmetry of peering partners (same size, geographical coverage etc). Other (more controversial) conditions are the exchange of roughly symmetric amount of traffic between the peering partners. This is reflected by "traffic ratios" which are the ratios between incoming and outgoing traffic.

Now, if the two ISPs are not symmetric in size, IP-transit arrangements come into account.
In IP-transit agreement, one provider ISP offers a customer ISP an Internet transit service, that is the global *reachability* to all the hosts connected to the Internet. As stated by BEREC *"the internet/broadband access provider pays for connectivity to the upstream network for upstream and downstream transmission traffic."* [1][14]. The customer ISP pays its provider for both upstream and downstream traffic. This service is charged based on volume consumption, regardless of whether the destinations that are served are close or far.

Furthermore, there are several levels (tiers) of networks of different sizes. Out-payments are made between these networks based on traffic speed and volumes [8][15]. The providers differ in size, market coverage and position in the market. Bigger Tier 1 providers will generally not peer with smaller Tier 2 or Tier 3 carriers. However, even if the peers are not symmetric, they might still decide to peer. In this case, there are financial settlements that are involved in the interconnection agreements, in order to compensate the network costs [9][16]. Such peering arrangements are not for free, they are called paid peering. Due to the fact that peering arrangements only come into account between symmetric interconnection partners with nearly symmetric amounts of exchanged and transported traffic, peering partners often disclaim billing the traffic.

There are several arguments against the current interconnection model. In fact, reachability, which is the traded good between ISPs, is vague; it is destination agnostic [10]. In transit for instance, the price paid to reach a farther (and therefore costly) destination is the same as that of reaching a destination which is near.

This vagueness creates an ecosystem that does not give incentives for ISPs to handle more traffic and invest in their network infrastructure. As a matter of fact, each ISP simply employs hot potato routing to get rid as soon as possible the traffic that gets into its network.

Another consequence is the "backbone free riding effect". If two ISPs that do not have the same size peer, then one of them, the smaller, free rides the bigger ISP. For example, ISP A, builds up and maintains a backbone network connection in different regions of a certain country, while a second ISP B just concentrates on a single region without building up a nationwide backbone. If these two ISPs agree on peering, carrier B uses A's backbone capacity to exchange traffic with A's customers in distant regions for free.

[11] cp. Kende, M. (2000), p. 5.
[12] cp. Norton, W. B. (2002), p. 4; cp. Bartholomew, S. (2000), pp. 33.
[13] p. 19f.
[14] p. 18.
[15] cp. Brenner, W. et al. (2007), pp. 51.
[16] cp. Davoyan, R. et al. (2009), p. 86.

Other aspects of the free riding effect can come in the case of a user-heavy ISP peering at *distant geographical locations* with a content-heavy ISP [10]. Due to the use of hot potato routing, the content-heavy ISP carries for a longer distance the requests for content (which is a small amount of data), while the user-heavy ISP carries (for a longer distance) the content itself which is a much higher amount of data. This situation results in one ISP carrying most of the traffic involved in the peering relationship and incurring therefore a much higher cost and a less benefit.

2.2 Recent Evolutions or is the current Internet still virtuous?

In [1] and [2], BEREC is postulating that in the current ecosystem, there is no free-riding and the ecosystem is virtuous:

- Both CAPs and end-users are paying for the Internet access
- Interconnection agreements are sufficient to compensate for inner network cost
- The decrease of equipment costs compensates for the traffic increase for operators

In our view, it is important to recall several facts which are insufficiently highlighted in the documents:

a) *Current quality and (past) over-provisioning*

The infrastructure of the current Internet has been funded by old telephone services, cable TV services and managed/specialized services which bring much more revenues than the pure so called data service of best-effort access to the Internet, typically priced at a flat-rate: for instance in 2010, while mobile voice and SMS services represented around 68% of the global mobile telecommunication revenues, mobile data services did not exceed 9%[17] .

In the current situation, best-effort is perceived with a sufficient quality. However, this quality results from past over-dimensioning by the ISPs to accommodate an increasing number of customers. As this number has stabilized recently (at least in developed countries), ISPs do not perceive the same incentives to invest. Furthermore, the migration of usages toward IP-supported services (in particular voice) make ISPs fear that the only - remaining service for them would be data transport. As a result, due to the low revenues generated by data transport services[18], network investments have been slowed down. ISPs are seeking for other source of revenues, and are trying to diversify their service offerings (through vertical integration for instance[19]).

b) Cost are not homogeneous for all actors

The Internet has been designed according to the end-to-end principle which imposes that the control is exercised by the hosts. The convergence of services at an IP level (e.g. voice, TV) has created, together with the end-to-end principle, a situation where OTTs stress the ISP infrastructure, while creating their own source of revenue. The OTT's revenue is not directly related to the ISP infrastructure. Consequently, the balance between ISP's investments and revenues is threatened.

BEREC's argumentation implies that (1) CAPs pay their share in the network costs (they pay to access the Internet), that (2) there is no free-riding effects on the current Internet and that (3) equipment cost reductions are compensating the increase of traffic volume.

On the first point (1), it is important to highlight that costs are very different depending on which network segment is considered. In the core of the Internet, where OTTs connect with transit ISPs, such costs are very low compared to the access segment. In

[17] NQ Logic, Broadband goes mobile, http://www.nqlogic.com/2010/04/broadband-goes-mobile.html

[18] NQ Logic, Broadband goes mobile, http://www.nqlogic.com/2010/04/broadband-goes-mobile.html

[19] Comcast purchsed Fardango, http://gigaom.com/2007/04/12/quietly-a-web-giant-takes-shape/, Orange took 49% of Dailymotion, http://gigaom.com/video/orange-dailymotion-deal/

Ireland for instance, looking at the usage-based (per Mbps 95th percentile measured at DSLAM) part of the price that the Irish incumbent charges[20]to alternative operators that do not own a part of the access infrastructure, one can see that the price per mbps per month in the transit[21] can be at least one order of magnitude lower than these prices per mbps in parts of the access network. Therefore, CAPs do not contribute that much to the network costs compared to the other actors (e.g. end-users and even business customers whose access bills are much bigger). In fact, in an Internet communication between a content provider and an end-user, the content provider pays at least ten times less than the end-user. This is due to the fact that in the current model, each actor pays only its access and not the entire Internet communication. Since content can "move", and users can not, content providers profit from the economy of scale to pay their Internet access much cheaper than end users for instance (they can settle where the prices are the lowest, and they aggregate a lot of servers). We conclude that the costs of an entire "Internet communication" (access-transit-access) are not equally shared between the actors that participate to this communication.

This unbalanced situation is driving a lot of research work on the topic of "fairer" sharing mechanisms among ISPs and OTTs. Among these works, Ma & al. [11] propose a revenue sharing mechanism in which each actor perceives compensations according to their contribution in the service. The work of Musacchio & al. [12] defines a two-sided market where ISPs receive a part of the advertisement profits generated by the OTTs. Some other works explore the possibility to create network services offered by ISPs that are richer than simple bundled reachability. Valancius & al. [13] for instance propose to rethink the transit product and argue that transit ISPs should differentiate their prices according to the destinations (closer destinations being cheaper than farther ones).

On the second point (2), we showed in Sec.2.1 that the nature of the interconnection between the Internet actors does not prevent from free riding effects, at least in theory. In fact, it is not possible to obtain a precisely fair peering relationship since the actors can never be "precisely" symmetric. Worse, it is even hard in practice to accurately estimate the benefits of a peering relationship for potential peers[14]. And even if a peering is balanced and fair at a given time, it can become fast unfair if the nature of the traffic of one of the peers changes (e.g. become content-heavy). We believe that the recurrent peering disputes are a telling symptom of the ill-suitability of the current interconnection model to prevent such effects.

On the third point (3), and even if the *"technical progress leads to cost improvements (on a per unit basis) […]"*[22], it may not be the only reason for the equipment cost reduction. Another hard competition between telco vendors evolved into a price war, allowing ISPs to buy equipment at lower-costs. Still, some infrastructure investments remain undone or may appear as blocked, which may affect the overall ICT economy at medium/long term. In Europe in 1Q12, investments fell 1% to $15.0bn [23]. The recent announcements of job cuts by telco vendors (ca. 20000 job cuts worldwide for Alcatel-Lucent, Cisco, NSN) confirm this trend and the low incentives for operators to invest in new infrastructures (either telco infrastructures or more applicative ones which are still marginal in volume[24]). Therefore, looking forward, further reduction in equipment costs cannot fully compensate for the loss of revenue experienced by ISPs, resulting in an inevitable reduction in infrastructure investments.

[20] Eircom's Bitstream Access Reference Offer (BARO) http://www.eircomwholesale.ie/Reference-Offers/Documents/Bitstream-Price-List-v7-13/ p.8

[21] http://drpeering.net/FAQ/Why-do-Internet-Transit-Prices-Drop.php

[22] [1], p.32

[23] http://ovum.com/2012/07/04/telecom-vendors-continue-to-bleed-despite-1q12-capex-rise/

[24] http://ovum.com/2012/07/04/telecom-vendors-continue-to-bleed-despite-1q12-capex-rise/

According to [15][25], returns on investments really differ among actors of the Internet Value chain. The conclusions show that, while each actor obviously has its own investments to develop its market, the profitability of the actors on different segments is different. Thus, except for the content right category which has a Return On Capital Employed (ROCE) of 14%, both head ends of the Internet value chain (Online service providers and User interface providers) have ROCE of 20+%. On the other side, enabling technology service providers and Connectivity providers have a culminating ROCE of 13%. These figures show that the past equilibrium which made the Internet "virtuous" may be broken, in particular for intermediaries of the Internet value chain. Regulators should therefore, as expressed previously, consider the differences in this global Internet value chain, and should allow all the actors to innovate: This should include the intermediaries who are the most constrained actors of the chain at the moment. Such innovations should trigger new business models that (1) preserve incentives to invest for any actor, (2) are compliant with regulators' aim to promote fair and transparent competition between the actors, and (3) aim at increasing social welfare for consumers.

Finally, the analysis is based on a single source of information[26]. We presented above another viewpoint that provides financial facts contradicting the conclusion obtained in this source. It seems for us therefore important to explain why BEREC has chosen such source, rather than basing its recommendations on multiple sources of information.

c) Effects of consolidation and integration in the Internet ecosystem

The aforementioned reduction of the scope of the neutrality debate seems not representative of the current market consolidation, with vertical and horizontal integration for most of actors of the Internet value chain. While the case of vertically Integrated operators is highlighted in [2] as potentially leading unfair treatments from operators to favour their own services, the effect of vertical integration for CAPs (e.g. Google, being at the same time the main search engine, a major content owner, a provider of caching and data transport services, and also the software editor of a major web browser and of a widespread operating system for handsets and tabs) or of horizontal integration for CDNs (not only providing caching services but often also transport services between main caches) are not really analyzed.

Vertical integration is currently applied by Network Service Providers (NSPs), i.e. bundling services on top of the Internet access including the access to music and video catalogues, telco VoIP, broadband TV programs in Standard or High quality definition. Such vertical integration is already covered by existing regulation. While risks may effectively exist for the Internet Neutrality considering these integrated operators, during the last decade in Europe, such triple play services had driven most of the broadband access deployments, and highly contributed to the current overall access quality to Internet and to the high penetration rate of Internet access services. All actors of the value chain are now benefiting from these improved penetration rates and service quality, including over-the-top service providers.

Generally, the documents insufficiently recall such positive effects in the current Internet ecosystem. [1] very quickly cites such possible positive effects, *"positive differentiations may be similar to specialized services, which do not necessarily raise competition problems, as long as they leave enough quality for the best effort delivery of traffic."*[27]*, or "On a longer timescale, beyond an adequate level of competition, a sufficient incentive to invest is needed for ISPs to foster the development of broadband infrastructures (that is next generation access networks). Differentiation practices, like charging end-users or CAPs for a better quality of service, may help operators to develop their revenues. Insofar as these additional earnings may contribute to the funding of networks (i.e. they*

[25] Chapter 5, p.17-18
[26] WIK consult reports, 2008 & 2011
[27] p. 58

correspond to reasonable and sustainable business models covering the costs of the infrastructure) they would have a positive effect on the long-term users' interest which have to be compared with other, possibly negative, effects."[28]. However, the conclusion of [1] poorly emphasizes such positive impacts (learning from the past), and insists on possible threats from differentiation.

It is worth noting that vertical and horizontal integration for other actors are introduced but not analyzed in the same way as vertical integration of ISPs. As an example, in [1], it is said that *"When a large CAP like Google operates its own networks and also peers this can be interpreted as a means of economizing expenditures for upstream capacity. It also contributes to increase user-experience by reducing latency as content is exchanged directly with another network instead of traversing one or more transit networks"*[29]. This explicitly recognized that a big CAP as Google provides its own transport services. It is however not discussed how such CAP transport services compete with the transport services of ISPs, and more importantly provide important market power to Google which competes with ISPs' transport services, or if Google is treating all content in a neutral manner (its own content and third party content cached or transported by Google).

In [3][30], it is mentioned that *"CDNs as a principle do not raise net neutrality issues, but discriminatory treatment in their favour might well do so, especially since large content and application providers and CDN providers are usually present at both the application layer (which is where the servers belong) and at the network layer. These aspects are outside the scope of these guidelines but are studied in the separate BEREC draft report on IP interconnection in the context of net neutrality."* This suggests CDNs do not only provide caching services but often consolidate their architecture with transport services. On the contrary, in [1] [31], it is said that *"CDNs do not interfere with the network layer of the ISPs. They do not provide connectivity but operates on top of the network layer on upper layers and in that sense can be qualified as a CAP"*. Latter, [1][32] however splits CDNs into two categories*: "[…] The CDNs runs servers and buys connectivity to the Internet for transmission between its servers like any other application provider. […] there are some reasons that the core functionality of a CDN may not be held to qualify as an ECS […] In an infrastructure based model, the CDN also runs the infrastructure to connect his servers and offers in addition to his core functionality to transport the CAP's data via this infrastructure. With regard to this offer for transmission services […] this service could qualify as an ECS"*.
While the classification seems relevant, the associated examples are a bit surprising. For instance, Akamai is mentioned as a pure server CDN, whereas Limelight is considered as having a network. Looking at some facts on peeringdb.com, one might observe that:

- The Akamai (AS20940) is present in 60 public IXPs and 35 private ones
- The Limelight (AS22822) is present in 41 public IXPs and 63 private ones.

In the ETICS project, a study on the Internet topology also revealed that Akamai has 32 ASes that are connected to 200 ASes (including Tiers 1 and other CDNs) and announces 1000 prefixes. Meanwhile Limelight has 12 ASes connected to 100 ASes and announces 45 prefixes.
Concerning the CAPs categorization, similar observations can be done as most of the big CAPs have a transport network. As an example, OVH (AS 16276) is present in 26 public IXPs and 16 private ones and Google (AS 15169) is present in 67 public IXPs and 70 private ones. These actors have their own network to interconnect their sites in all these points, act somewhat as network operators, but they're not considered (nor regulated) as such.

[28] p. 29
[29] p. 38
[30] p. 19
[31] p. 14
[32] p. 16-17

In [1][33], it is concluded that *"The emergence of hyper giants (Google, Akamai, Amazon…) as well as new kinds of peering arrangements (e.g. regional peering) significantly contributed to the flattening of the Internet topology. Boundaries between players in the value chain becoming more blurred as players increasingly perform different and/or new functionalities and integrate along the value chain"*.

Taking into account such an evolution of the Internet ecosystem and in the context of a wider debate on the Internet Neutrality, we recommend to not oppose and have a similar treatment between integrated CAP-driven, CDN-driven or ISP-driven possible innovations to improve the future Internet sustainability for all actors, as well as social welfare for Internet' users.

3. What level of QoS? What are current and future needs for QoS?

3.1 What level of Qos?

Three categories of services could be differentiated according to their QoS requirements:

The first category of services is characterized by no special QoS requirements (e.g. E-Mail, web browsing). This means that these services are not quality-sensitive. Therefore the best effort manner of the Internet is sufficient for the realisation of this kind of services, which means that all kinds of traffic are always treated equally and independently of the service-specific QoS requirements.

Compared to the first category, services of the second category are highly QoS-sensitive (e.g. Cloud computing, quality-sensitive online-gaming, e-Health). Therefore today's best effort IP-transport in the Internet is not sufficient for their realisation. The equal treatment of every bit transported via the internet would lead, particularly in terms of congestion, to negative interference of this kind of services. If high quality transport is required for this category of services, the application layer may be built so as to make them tolerant enough to run properly with a quality not as guaranteed as in the third service category. However, the Internet's value is about services not bits. Since there are delay-sensitive and delay-tolerant applications by definition, care should be taken that both these categories are properly managed by the Network Layer, by identifying and meeting their different needs.

Indeed the QoS of this second kind of services is realised best by the implementation of different QoS-classes on the transport level. This means that each traffic flow is mapped to a special class of quality that is adequate to its requirements (Bandwidth, jitter, delay, packet loss…). With the IP's typical separation of the service/application level from the network/transport level, different QoS-classes on the transport level could enable "relative" QoS "within the public internet". This means that the quality could be associated to an IP traffic flow shared by many applications and services requiring the same quality objectives at the IP transport level.

Finally, the third category of services like high quality voice or IP-TV should be provided with a guaranteed QoS to the customers as is it is done for today's PSTN or TV networks. Therefore, these services may be realised "outside the public Internet" on dedicated networks. In the case of high quality voice, it may be realised with dedicated interconnections between the networks to guarantee both the voice-specific quality and features, as also stated in the BEREC consultation documents. At the points of interconnection, session border controllers with dedicated interconnection lines between them are implemented.

[33] p. 48

In contrast to the other two service categories, in this third case, an additional control plane builds a link between the application and the network layer to control and manage the service realisation and to ensure and guarantee the service-specific quality as well as service-specific features. Furthermore the QoS of these services should not be interfered with other services.

In the above examples, we mainly discussed how classical QoS parameters (availability, bandwidth, loss, delay, and jitter) need to be more or less strictly assured. Additional performance criteria can be associated to the network service, such as specific routing and/or security preferences. Indeed, for certain critical services, a customer may ask for guarantees and constraints on the route its data follows between the two extremities, for example to ensure that its data do not cross a given country for security, legal or regulatory reasons. Such performance criteria may impose additional interconnection processes and business models.

3.2 What are current and future needs for QoS?

Due to the fact that IP-based networks are multiservice networks, the different service-specific requirements should be taken into account. As described in the BEREC consultation document, today's internet is based on a "best-effort" principle[34], which means that every service transported via IP-based networks is treated equally and independently of its quality of service requirements. The trouble is that they are all treated equal to "no guarantee"; which means it can happen to work properly (this was the case of the first generation of over-dimensioned networks within the Internet), and it can equally happen to be congested. Another major issue is that congestion is not exposed and the cause of it cannot be policed. Instead, all users and flows suffer from "aggressive" user flows and opportunistic behaviour that harms social welfare [16].

This leads to problems in the realisation of services particularly in case of congestion. Quality-sensitive services could only be provided with a poor QoS or would even have to be crowded out. In order to avoid congestion and the crowding out of quality-sensitive applications and services, capacity extension needs to be complemented with sensible and customer-oriented network management.

The need for QoS is not presented in a consistent way along the documents. For example, explanations are provided to argue that the Internet provides a sufficiently good quality and that no traffic differentiation is required:

- In [1][35], *"Best-effort Internet in most cases results in a (relatively) high quality of experience for users, even for delay-sensitive applications such as VoIP".* VoIP is a real-time service with certain needs in terms of end-to-end delay and jitter. But innovation at the application layer from VoIP providers such as Skype indeed allowed a reasonable quality for voice communication over the Internet, making it more robust to delay variations and packet loss. However, Voice is also a service consuming a low bandwidth and thus do not impose strong constraints and congestion in ISPs' networks. Drawing general conclusions from this specific service seems not appropriate and may be detrimental for future advanced services with more stringent constraints.
- Also in [1][36], *"Examples are endpoint-based congestion control for reduction of the traffic load, Internet Exchange Points and increased use of peering. Also CDNs are used to improve the CAU's perception of an application's quality (QoE).".* Congestion marking must be done at intra-domain level so as to be efficient. This does not happen now, so huge amounts of traffic can be transferred from USA to Europe just to be dumped in the last mile. Caching by CDNs indeed provide greater users'

---

[34] See BEREC (2012), An assessment of IP-interconnection in the context of Net Neutrality, Draft report for public consultation, 29 May 2012, p. 26

[35] p. 28

[36] p. 28

experience, but is not applicable to all kinds of services, in particular those requiring person-to-person live interactions. Even more, as explained in [5], the combination of CDN-"neutral" best-effort paradigm clearly favours certain applications in the expense of others (highly QoS-sensitive) thus hindering the latter's deployment.

Quoting [1][37], *"CDNs are used to improve the CAU's perception of an application's quality (QoE)."* and *"The market has developed very well so far without any significant regulatory intervention"*[38]. CDNs' existence is somewhat a sign of inefficiency: they also cause some kind of differentiation since CAPs having a CDN contract have de facto a preferential treatment of their content as opposed to those who do not. In addition, there is no transparency on the contract from CDNs and no data are available to know if all clients receive the same treatment in the CDN infrastructure (whether it concerns the access to CDN caches or also includes a data transport service). These questions have not raised strong concerns for the moment, nor reduced the capacity of innovation on the Internet, even if not all CAPs can afford to pay for such a premium service.

It is therefore demonstrated from the OTTs' innovations that quality improvement is required for certain applications due to the unpredictable behaviour of the network. Drawing general conclusions from specific current services, for which stop-gap solutions exist, to argue that QoS differentiation in the network is not required in the future is not appropriate and may be detrimental for future advanced services with more stringent constraints.

In the near future, more QoS will be required: as an example, cloud services already represent a significant market for enterprises and end-users (Gardner estimated the SaaS revenues to $10.0 billion in 2010 with a growth projection to $21.3 billion by 2015[39] and additional growth projection to $10.5 billion for IaaS[40]). Despite these huge markets, a major limitation is the absence of Service Level Agreements for the network part in cloud offerings. The situation introduced by Lonsethagen & al. [17] in 2011, is still present in 2012 [18] with a significant difference: the appearance of cloud offers from operators introducing such differentiator. While being well positioned to offer such integrated features, operators could also provide in the future such assured network service on-demand for over-the-top cloud service providers, thereby increasing the competition and social welfare. The increase of energy price in transports, as well as a general greening effect of the society, will also increase the needs for professional video-conferencing and telepresence services, having real-time and bandwidth constraints.

4. Positive and negative impacts from QoS management

4.1 QoS as a new service

The implementation of QoS-mechanisms on the IP-transport level on top of today's best effort internet should not change the interconnection regime for the best effort class. Peering and transit arrangements could still be maintained as a charging mechanism.

In contrast, the interconnection regime for the QoS-mechanisms on top of best effort should reflect the provision of higher end-to-end QoS delivery and the greater value for the custumers and the OTTs, as compared to best effort. Therefore it is rational and efficient that the charging mechanism for the QoS-classes leads to higher prices than for the best effort

---

[37] p. 49
[38] p. 50
[39] http://www.logisticsmgmt.com/article/gartner_says_worldwide_saas_revenue_is_forecast_to_grow_21_percent_in_2011/
[40] http://www.qas.com/company/data-quality-news/iaas_market_to_record_strong_growth__7178.htm

Internet-class. Moreover this would also increase the incentive to further invest in new network equipments and systems allowing the support of QoS-mechanisms.

This rational is analogous to service differentiation in other domains. For example, due to increased usage of cars for commuting in urban areas, the transportation system became congested. One solution for that was to add toll highways, by which drivers that are willing to pay a premium, enjoy a better service (less congested highways, shorter commute time). As a side effect, traffic volumes on the best effort infrastructure (the toll-free highways) are less stressed. The QoS service classes could be seen as premium services that have strict traffic demand (bandwidth, delay, jitter), for customer who are willing to pay for them. This rational is productive, as long as the best-effort service is still valuable (namely, the freeways are still usable). We believe that this should be the focus of the debate.

In light of this example, BEREC's response statement that *"all electronic communication passing through a network is treated equally"* should be questioned. Should all cars in our roads treated equally? Why is that trucks are excluded from certain roads or lanes for some specific time of the day? Why do we have carpool lanes or dedicated lanes for public transportation? Clearly, in matured markets, the product strategy of "one size fits all" is not appropriate. Reachability is no longer the main objective of a mature transportation system; rather, the efficiency and the service level of the transportation system should be considered. The state and the uptake levels of internet services suggest that multiple service levels should be made available, accommodating different needs, budget and context, improving efficiency, while contributing to an improved social welfare.

Another example for service differentiation is the mobile telephony service, which can be seen as a premium telephone service. Clearly, mobile telephony services are priced much higher than PSTN services that are justified by the premium service that is delivered to users (mobility). The market evolves into a situation where consumers have both, mobile and fixed telephony services. We should see the same practices in the Internet world, where best-effort and premium QoS services are available to consumers, and they use the service that is most appropriate to their budget, context and application.

ETICS proposed solutions for a new interconnection framework. The goal of this contribution is not to describe in details research activities and proposed solutions of the ETICS project. However, we found relevant in the context of the BEREC consultation to show possible evolutions that may arise from Inter-Operator QoS management. These evolutions are expected to create a more sustainable role for the different actors in the Internet value chain. **Such proposals have to be taken as illustrative example and other solutions allowing the same or partial features could obviously also be envisaged.**

The main concept is to allow Network Service Providers to exchange, new types of goods, between them (i.e. between NSPs) or between NSPs and their customers (CAPs, OTTs, end-users). Indeed, as we previously indicated, we think that the traditional reachability service has reached its limits, and that a new type of service should complement such basic interconnection service. We thus defined ***Assured Service Quality paths (ASQ paths)*** as quantifiable network connectivity services that guarantee quality data carriage (e.g. bandwidth, delay, etc.).

Today, it is possible for network operators to offer assured quality network connectivity services, if such services are provided via a network infrastructure managed by a NSP. Examples of such services include business connectivity services where VPN operators interconnect the different sites of their business customers. One can even consider the example of triple play-managed services in which access providers give guaranteed network connectivity between their users and their TV or VoD platforms.
However, making such assured quality network services cross the boundaries of single domains remains elusive, even after almost 20 years of research on the topic. **ETICS believes that enabling such inter-carrier assured quality connectivity services**

requires both, technical solutions and an economical ecosystem with the right mechanisms for operators to trade such new network products. To this aim, ETICS introduces a new plane, the network service and business plane. This plane lies on the top of the historical management, control and data planes. It allows NSPs to collaborate in order to offer assured quality network connectivity services that cross their domains. It is beyond the goal of the document to describe how this can be realized. For more details, please see the annex bellow and [19].

Some level of collaboration among NSPs is needed for offering end-to-end ASQ paths. One of our research [20], demonstrates that via such collaboration, all actors enjoy additional benefits (NSP, end user), while the social welfare is improved. [20] analytically proves that when NSPs have the same objective (e.g. minimization of end to end packet delay for a given bandwidth, a reasonable objective for interactive video services), via negotiation and binding agreement among all involved NSPs, the social optimum can be reached, while each NSP improves its utility. Even without all-hand common objective, negotiation and agreement between some of the NSPs improves the outcome of the worst-case possible scenario for each of the collaborating NSPs.

The "Sending Party's Network Pays" (SPNP)-principle that was investigated within ETICS appears to be a promising charging model for the QoS-classes on top of the best effort Internet. The SPNP-regime can be considered to be a fair alternative to current peering agreements that may be asymmetric due to unbalanced traffic flows, since both parties pay for the amount of traffic they send to the other network. Each party only pays for the traffic it generates which is economically efficient.

The SPNP model is inspired by similar solutions in other industries. For example physical packet delivery and mail services also hand over responsibility for delivery along with compensation, generally according to the quality or speed of delivery. Thus, the SPNP revenue model is expected to create general awareness about the fundamental cost of network traffic and motivate an efficient usage of bandwidth resources.

4.2 Divergent literature on the impact of QoS management

a) Allowing QoS management will degrade the Best Effort Internet Quality?

Traffic management techniques are generally perceived as providing possible positive and negative effects, which is also acknowledged by BEREC. Citing the conclusion of [1][41]:
" *p) Generally, QoS classes at the network layer can be welfare enhancing as long as CAUs can make an informed decision. This requires transparency.*
q) On the other hand, the introduction of traffic classes using prioritisation introduces an incentive to degrade the best effort class in an anti-competitive manner, in order to induce customers to pay the higher price for the managed traffic class.
r) Up to now, interconnection with QoS assured across network boundaries does not / hardly exist in practice."

The argument that introducing QoS will provide incentives for operators to degrade the Best Effort access is not demonstrated. Actually, since the early days of the net neutrality debates, opponents and proponents of QoS services have faced with various arguments. In [4] and [21], the authors consider that service differentiation can harm the end-to-end connectivity of the Internet and the innovation. Meanwhile, Shetty & al. [22] develop a model of QoS that provides social welfare under a simple regulatory tool. In [10], Ma & al. also demonstrate that the mechanism they propose has some desirable fairness properties when sharing the revues ISPs would generate with service differentiation. Jaminson & Hauge [23] advocate that premium services can also be beneficial for OTTs, fostering innovations at the edges.

[41] p. 49

If, as illustrated by the point/counterpoint of van Schewick & Farber [24], the literature can be considered as inconclusive about the benefits from QoS or service differentiation, all authors reckon or implicitly based their work on the assumption that expenses in the Internet are unbalanced.

More fundamentally, most authors acknowledge that today, de-ossifying the Internet economics is a prerequisite to any technical or service innovation thus meaning that the traditional reachability service has attained its limits in this field.

The impacts of QoS-mechanism on the Internet access charging can be considered as unpredictable since the charging of QoS could take various forms. But, we believe that, because the Internet access is a highly competitive segment market, failing in the performance challenge will result in losing customers for operators. Such operators would then immediately be penalised by the market. Furthermore, investments that will be made to offer high-quality service are very likely to benefit to all types of services (per analogy with the car industry, where high-end innovation initially made for high-quality cars, e.g. on breaks or security devices, has finally benefited to all categories of cars).

We do not think that this imposes a regulation on the minimum level of QoS to be guaranteed for the Best Effort Internet access. Moreover, the definition of absolute values minimum QoS parameters can be quite hazardous since the quality experienced by end-users varies according to applications. Meanwhile, such absolute values should evolve along with the technological evolutions of applications. Hence, we think that a public *ex-post* market survey from national regulators may be sufficient to detect any abuse from network operators.

b) Allowing QoS management will reduce the innovation on the Internet by raising the entry barrier?

As explained previously, a similar situation already exists with CAPs who can afford to pay the service of a CDN to improve the quality of experience of users of their applications services. This could be perceived at a higher entry barrier for the "man in the garage" to propose new services over the Internet. However, it is now proven that the introduction of CDNs did not reduce the capacity of innovation over the Internet, and, on the contrary such innovations even happened.

QoS services from network operators will be proposed on top of the Best Effort Internet, and, as mentioned earlier, market mechanisms already assure that operators do not have incentives to degrade the "Best Effort Quality". Consequently, it is not expected that the introduction of QoS management services will effectively increase the entry barrier for innovators. On the contrary, allowing for product differentiation will increase the variety of the offered products and services. This by itself is assumed to increase total welfare, as demonstrated in [20]. Finally, the introduction of QoS services would also provide incentives for application designers and developers to reduce the load of their applications on the network.

5. Conclusions

Reading BEREC's drafts, we, as experts involved in the technical and business network research area, have been surprised by the relative strong positioning of BEREC against QoS Management. Even if no strict regulatory recommendations were given, we would have expected from BEREC a more neutral analysis, highlighting negative as well as positive effects of such possible innovations. As such, we suggest that the opinions developed along the current version of the document may not affect all the actors of the Internet value chain in the same manner. In particular, this may be very detrimental to some European actors such as Network operators and equipment vendors, while those are already suffering from the crisis affecting particularly the Telecom Industry in Europe.

Based on a quite large literature review, we demonstrate from the pros and cons analysis, that no strong conclusions can be presently derived. In this answer, we thus tried to counterweight BEREC's analysis of the current Internet ecosystem. We believe that research and innovations on QoS differentiation should be considered and evaluated by the market to really allow each actor of the value chain to innovate. We also think that the deployment of such services is an opportunity to improve both the users' welfare as well as the economic sustainability of all the value-chain actors. Once such QoS-mechanisms are about to be deployed, the potential drifts accompanying them could be anticipated and ultimately regulated if necessary.

In general we strongly support the BEREC statement in the consultation document that "any [regulatory] measure could potentially be harmful, so that it should be carefully considered"[42] and we agree that the best-effort Internet works very well without regulation. In spite of the implementation of QoS-mechanisms on top of the best-effort Internet, the best effort Internet will remain as it is today. In fact, the Internet today works well mostly due to innovation and repurposing. And the absence of regulation has clearly allowed the Internet technology to be used in ways that were never imagined at the net's creation era, leading to what we perceive today as "The Internet", an ecosystem which is constantly evolving and enriched with new services. This is a valuable lesson that posing limitations and forbidding innovation may hinder positive effects that are unforeseen even by top specialists in the present era.

Furthermore there is no harm of the Net Neutrality-principle by implementing QoS mechanisms on top of the best-effort Internet to ensure end-to-end QoS-connectivity. Every service will be treated equal in each QoS class and every customer and OTT has the free choice in what manner their traffic should be transferred. Additionally, they could still use the best effort Internet in the same manner as today.

By endorsing the concept of "quality-based delivery", it will be possible to establish new interconnection policies based on the "value" of the traffic (not only on the "volume"), enabling new business models and implementing an ecosystem where operators' revenues will not be disconnected from the investment needed which would be necessary by the rapid growth of Internet traffic.

The availability of differentiated products at competitive prices will further stimulate the digital economy and ultimately everyone would benefit. Consequently competition in the provision of QoS-levels would manifest in the form of price competition and in quality competition.

In sum, the introduction of end-to-end quality will increase the variety of the offers and the customer choice. If such quality is openly and transparently applied, we can expect them to improve the global welfare and also foster competition and innovation.

---

42 [1], p. 50

References:

[1] BEREC, "An assessment of IP-interconnection in the context of Net Neutrality - Draft report for public consultation", 29 May 2012.

[2] "Differentiation practices and related competition issues in the scope of Net Neutrality - Draft report for public consultation", 29 May 2012.

[3] "Guidelines for Quality of Service in the scope of Net Neutrality - Draft for public consultation", 29 May 2012.

[4] Wu, Tim, "Network Neutrality, Broadband Discrimination". In Journal of Telecommunications and High Technology Law, Vol. 2, p. 141, 2003. Available at SSRN: http://ssrn.com/abstract=388863 or http://dx.doi.org/10.2139/ssrn.388863.

[5] P Faratin, David D. Clark, Steven Bauer, William Lehr, Patrick W. Gilmore and Arthur Berger, "The Growing Complexity of Internet Interconnection", Communications & strategies, No 72, 4th Q. 2008, p.51-71.

[6] Kende, M. (2000): The Digital Handshake: Connecting Internet Backbones, FCC Office of Plans and Policy, Working Paper Number 32, 2000, Retrieved March 26, 2011, from http://www.fcc.gov/Bureaus/OPP/working_papers/oppwp32.pdf

[7] Norton, W. B. (2002): A Business Case for ISP Peering, Retrieved April 04, 2011, from, http://pages.cs.wisc.edu/~akella/CS640/F06/readings/norton.pdf

[8] Brenner, W.; Zarnekow, J.; Kruse, J.; Sidler, A. (2007): Qualität im Internet – Technische und wirtschaftliche Entwicklungsperspektiven, St. Gallen, 2007.

[9] Davoyan, R.; Altmann, J.; Effelsberg, W. (2009): A New Bilateral Arrangement between Interconnected Providers, in: Reichel, P.; Stiller, B.; Tuffin, B. (Eds.); Network Economics for Next Generation Networks, Berlin, 2009, pp. 85 – 96.

[10] Z. Ben Houidi and H. Pouyllau "The price of tussles: bankrupt in cyberspace?", In Proc. of ACM SIGMETRICS/performance Workshop on Pricing and Incentives in Networks 2012.

[11] R. T. B. Ma, D.-M. Chiu, J. C. S. Lui, V. Misra and D. Rubenstein, On cooperative settlement between content, transit, and eyeball internet service providers, IEEE/ACM Transactions On Networking, vol. 13, no. 3, pp 802-815, 2011.

[12] J. Musacchio, G. Schwartz, and J. Walrand, "A Two-sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue," Review of Network Economics, vol. 8, no. 1, March 2009, pp. 22-39.

[13] V. Valancius, C. Lumezanu, N. Feamster, R. Johari and V. V. Vazirani, How Many Tiers? Pricing in the Internet Transit Market, ACM SIGCOMM, pp 347-356, 2011.

[14] A. Dhamdhere, C. Dovrolis, and P. Francois, "A value-based framework for internet peering agreements," in International Teletraffic Congress (ITC 22), 2010.

[15] A.T. Kearney, "Internet Value Chain Economics", January 2011.

[16] T. Moncaster, J. Leslie, B. Briscoe, R. Woundy, D. McDysan, "ConEx Concepts and Use Cases", draft-ietf-conex-concepts-uses-01, March 2011.

[17] Håkon Lonsethagen et al. "Cloud connectivity, interconnect and SLAs - From a network perspective", Spring Future Internet Assembly, Session III.1, "The Network Lost in the cloud", Budapest, May 2011.

[18] João Soares, "Business Aspects of Cloud Networking", Future Network and Mobile Summit, workshop on Cloud Networking – Technical and Business Challenges, July 2012.

[19] Antonio Cimmino, Luigi Damis, Richard Douville, Zied Ben Houidi, Nicolas Le Sauze, Olivier Dugeon, Eleni Agiatzidou, Finn Tore Johansen, "Network architectures for end-to-end business and traffic collaborations among carriers", Future Networks and Mobile Summit conference, July 2012.

[20] Gideon Blocq and Ariel Orda, How Good is Bargained Routing, Infocom 2012, Shanghai, China, April 2012

[21] J. Crowcroft, Net Neutrality: The Technical Side of the Debate, ACM SIGCOMM Comput. Commun. Rev., vol 37, no , pp. 49-56, 2007.

[22] N. Shetty, G. Schwartz, J. Walrand, Internet QoS and Regulations, IEEE/ACM Transactions On Networking, vol. 18, no. 6, pp 1725-1737, 2010.

[23] Jamison, Mark A. and Hauge, Janice Alane, Getting What You Pay for: Analyzing the Net Neutrality Debate (April 20, 2008). TPRC 2007. Available at SSRN: http://ssrn.com/abstract=1081690 or http://dx.doi.org/10.2139/ssrn.1081690

[24] Barbara van Schewick and David Farber. 2009. Point/Counterpoint: Network neutrality nuances. Commun. ACM 52, 2 (February 2009), 31-37. DOI=10.1145/1461928.1461942 http://doi.acm.org/10.1145/1461928.1461942

Annex - ETICS proposed solutions for a new interconnection framework

It is not the goal of this contribution to discuss the technical details, merit or possible technical limitations of the proposed ETICS collaboration schemes. Below, we just quickly introduced main concepts as examples of possible innovations for future QoS-enhanced interconnections.
Details can be found in [17], in public deliverables available on the project web site[43], or by contacting the authors.

A) New Assured Service Quality (ASQ) products and services
An ASQ product represents a connectivity service with defined guarantees (soft or hard) for QoS parameters such as the availability, the bandwidth, the delay, the jitter, etc. For example, Fig. 1 illustrates two basic ASQ services:

- An ASQ-Traffic termination (ASQ-TT): the customer NSP wants to reach users in a destination region within the provider NSP domain, with certain QoS attributes. The ASQ-TT product consists in the ASQ traffic delivery of traffic from the point of interconnect (PoI), to a destination region.
- An ASQ transit: it corresponds to the traffic delivery through a transit NSP from one traffic delivery source point (TDSP), to a traffic delivery destination point (TDDP).

ASQ-TT TO A DESTINATION REGION (1)	ASQ TRANSIT (2)
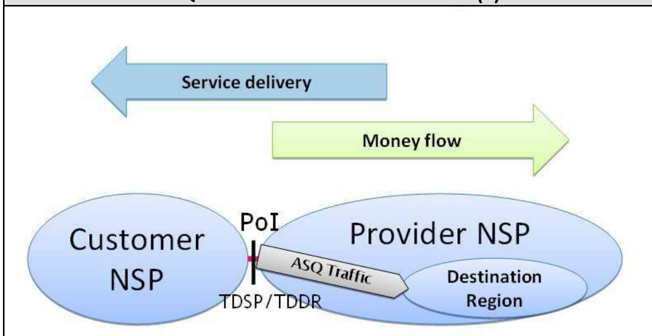	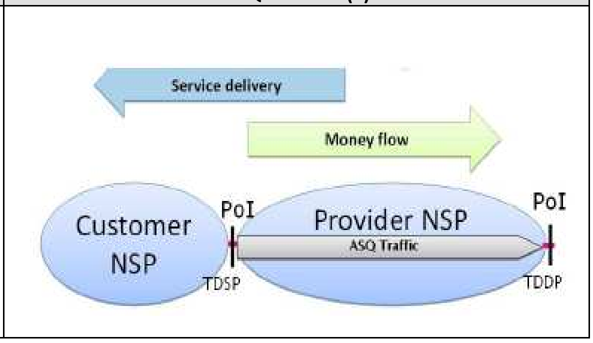

Fig. 1: typical ASQ products offered by a provider NSP to a customer NSP.

It can be observed from the figure 1, that the money flow compensates the new service sold by the Provider NSP. In an ETICS context, each NSP is remunerated according to the network service it provides.

B) Discovery and composition of per-NSP ASQ products
Without describing the ETICS reference architecture, we explore in the following sub-section how communities of customer and provider NSPs can cooperate to compose ASQ products in order to satisfy a customer demand. The composition of multi-NPS products will include two tasks:

- Service Discovery - to find the right set of NSPs that can provide the multi-NSP ASQ product (an ASQ product spanning multiple NSPs);
- Service Composition - to allocate individual QoS objective per NSP.

There are actually multiple probable scenarios to compose multi-NSP products, including cases where the service composition and discovery tasks can be done. Key aspects to consider are:

- Separately (on-demand product offers composition or "**pull mode**"): the composition of NSP products starts only upon the reception of a customer request
- At the same time (pre-computed product offers composition or "**push mode**"): NSPs products are pre-computed in the form of well-defined multi-NSP offers. This type of scenarios is called push model because NSPs push their offers (of ASQ products) into a repository, also known as service catalogue. These offers need to be shared by

---
[43] www.ict-etics.eu

other members in the community in order to be able to compose and build a final product

- Inter-NSP Information exchange scenarios for respectively the PoI-to-Region, and PoI-to-PoI services.

Then, it is important to understand how such products can affect the possible interconnection mechanisms and business models between NSPs, as well as business models among NSPs and CAPs. We actually analyzed different collaboration schemes, which both exhibit different technical and business properties. These schemes range from fully distributed to fully centralized, as well as intermediate hybrid schemes. The precise specifications of these schemes are being finalized now in the last year of the project.