BEREC's comments on the ETNO proposal for ITU/WCIT or similar initiatives along these lines

In summary

The following are BEREC's comments on the ETNO proposals for the World Conference on International Telecommunications 2012 (WCIT 2012) and any similar proposals that might be made along those lines. They are based on the content of existing BEREC reports on IP-interconnection and Net Neutrality.[1]

The internet has developed well without regulatory intervention, through stakeholders' coordination in the free market. Its ability to evolve over time and self-adapt has been key to its growth and success. BEREC believes that the nature of services to be delivered across the network, and the charging mechanisms applied to them, should continue to be left to commercial negotiations among stakeholders. In such a fast-moving and dynamic field, to raise their profile through explicit reference in an international treaty revised every 25 years is inappropriate and could be prejudicial to the continued development of the Internet, having real potential to harm consumers and content and applications providers (CAPs), disrupt access to content, and contribute to a widening of the digital divide.

ETNO has proposed to include an explicit reference in the International Telecommunications Regulations (ITRs) to a specific interconnection charging mechanism (Sending Party Network Pays, or SPNP) and to the concept of end-to-end QoS delivery (operating in parallel to a best efforts Internet)[2]. The proposal is presented as simply highlighting and raising the profile of two approaches to service delivery which telecoms operators intend to promote in commercial negotiations. Separately, ETNO has argued that its textual amendments to the ITRs are simply aimed at ensuring that existing ("paid" peering) commercial arrangements do not come to be prohibited by law in any ITU member state.

However, BEREC believes that the elevation of these approaches (however ETNO may label them) by explicit reference in a (high-level) international treaty runs the real risk shifting the balance of negotiating leverage between market participants and inducing an abuse of market power by telecoms carriers in relation to terminating traffic (much as occurred historically in traditional telephony). The resulting shifts in market power would increase the need for regulatory oversight and potentially require regulatory intervention, in accordance with the EU regulatory framework.

Strictly speaking ETNO is advocating an "interconnection philosophy" based on transmission services being provided across the Internet all along a defined path between endpoints, much like the connection-oriented circuit switched "old generation" PSTN networks and voice services on which ETNO members built their businesses. This is fundamentally at odds with the principles of connection-less packet switched networks underlying the success of the Internet to date, based on decentralisation and simplicity. BEREC believes that the benefits of a connection-less network risk being unravelled by the widespread adoption of connection-based practices on the global Internet.

It is in all our interests to protect the continued development of the open, dynamic and global platform that the Internet provides, which has evolved over time (without regulatory intervention), and helped enable so much innovation at the network endpoints. ETNO's proposal could undermine this and therefore lead to an overall loss of welfare.

Specific BEREC observations on quality and charging

On Quality of Service

Fundamentally, there is no end-to-end awareness and dedicated transmission path on the (connection-less) Internet, as routing decisions are taken locally by autonomous networks, and packets included in one data flow may take different routes over separate networks. Interconnection at the network layer does not rely on application-based end-to-end information thanks to the separation of the application layer and the network layer that is inherent to the Internet.

While mechanisms for introducing differentiated QoS traffic classes have been available for more than a decade, BEREC notes that these have not been implemented across networks on the Internet (as opposed to the provision of specialised services within operators' own networks, e.g. in relation to IPTV).

This lack of demand may be the result of any one of the following reasons:
- While not providing a guaranteed quality level of data delivery, the best efforts approach of the Internet does not imply low performance, and in fact results in most cases in a high quality of experience for users, even for delay-sensitive applications such as VoIP.
- Customers are not willing to pay much of a premium for a better service, because they consider that their service is good enough.
- End-to-end Service Level Agreements are costly and unwieldy to implement because QoS has to be assured on each network along the path, end-to-end.
- Indeed it would be difficult to define what constitutes an end-to-end "premium quality" level and to assess whether it has been achieved in such a decentralised ecosystem.

Over the Internet, a guaranteed end-to-end QoS offer is, therefore, neither commercially nor technically realistic. Differentiated services (DiffServ), which fall just short of guaranteed end-to-end QoS, exist but continue to be exceptional, for the reasons listed above (and not because they are anywhere prohibited).

In the few instances where end-to-end QoS arrangements are currently in use, they almost always consist of specialised services (e.g. IPTV), provided not over the Internet but within a closed network within the Internet Access Provider's own network. This is then not a case of interconnection between networks and therefore not dealt with in this statement.

It is important to note that mechanisms other than end-to-end QoS traffic classes have been developed over time for improving end-to-end network performance, including end-point based congestion control for reduction of the traffic load, Internet Exchange Points and the increased use of peering. Content Delivery Networks (CDNs) are also used to improve the user's experience of an application's quality (QoE). All of these mechanisms have evolved through commercial innovation, without the need for regulatory intervention. Furthermore, they do not threaten the system of decentralised efficient routing of Internet traffic, since they are applied at endpoints. On the other hand, that system would be threatened by the

widespread implementation of traffic classes across the internet that have to be implemented *within* the network and also require an additional management plane on top of the IP layer.

While QoS differentiation may be an appropriate tool to deal with scarcity of bandwidth in access networks by prioritising, e.g., voice services, the situation is different in IP-backbone networks, where additional capacity is relatively cheap[3]. Considering the above, the only reason why ETNO has put end-to-end QoS on the table seems to be "to enable incremental revenues by end-to-end QoS pricing and content value pricing" as ETNO themselves have put it[4]. Operators have less control over transmission networks on the Internet than they used to have with voice traffic. Put simply, ETNO is trying to extract additional revenues from its existing network assets, in a bid to reassert control over a changing communications ecosystem.

On charging principles

BEREC has for a long time been looking at charging principles in the context of IP networks[5]. In the current generation PSTN, a "calling network party pays" approach is used at the wholesale level to charge the calling party (which initiates the transmission) for voice calls using a dedicated end-to-end connection in circuit-switched networks. Because of the termination monopoly of the telephony access provider, the terminating charges are regulated.

In contrast, interconnection on the Internet has operated on the basis of transit/peering arrangements at the higher level, and a "bill & keep" approach where the terminating access network operator does not receive payments at the wholesale level for terminating the traffic, but recovers its costs at the retail level from the end-user. If "bill & keep" were to be replaced by SPNP then the ISP providing access could exploit the physical bottleneck for traffic exchange and derive monopoly profits, requiring regulatory intervention.

IP interconnection agreements only involve the provision of capacity of the interconnection link and not the end-to-end transmission of particular data flows across different autonomous IP networks. Unlike voice traffic on old PSTN networks, data does not travel over an exclusive, dedicated network connection, and it is not possible to ascertain the nature or volume of a particular data flow end-to-end (and so not possible to charge for it that way either).

Therefore, charging for IP-interconnection usually takes place on the basis of the capacity provided at the interconnection point. ETNO's proposed end-to-end SPNP approach to data transmission is totally antagonistic to the decentralised efficient routing approach to data transmission of the Internet. The connection-oriented nature of end-to-end SPNP, with its focus on charging based on the actual volumes or value of the traffic, would represent a dramatic change from the existing charging framework operating on the Internet.

Furthermore, if other charging practices became widespread which enabled Internet Access Providers connecting end-users to set abusive charges for interconnection out of a monopoly position, this situation would need to be addressed.

Finally, it is worth pointing out that ETNO's proposals do not seem to have taken account of the fact that the request for the data flow usually stems not from the CAP but from the retail Internet access provider's own customer (who "pulls" content provided by the CAPs, and from whom the ISP is already deriving revenues). Ultimately, it is the success of the CAPs (from whom ETNO wishes to extract additional revenues) which lies at the heart of the recent increases in demand for broadband access (i.e. for the ISPs' very own access services).

Indeed, both sides of the market – CAPs on the one hand and users of these applications on the other hand – already contribute to paying for Internet connectivity[6]. There is no evidence that operators' network costs are already not fully covered and paid for in the Internet value chain (from CAPs at one end, to the end users, at the other).

This model has enabled a high level of innovation, growth in Internet connectivity, and the development of a vast array of content and applications, to the ultimate benefit of the end user. Attempts to undermine it could put these benefits at risk.

14 November 2012

---

[1] See www.berec.eu. ERG (07) 09 , ERG (08) 26, BoR (10) 24 Rev 1, BoR (10) 42, BoR (11) 53, BoR (12) 31, BoR (12) 32, BoR (12) 33

[2] More specifically ETNO suggested the following additions to the ITRs:

3.1 Member States shall facilitate the development of international IP interconnections providing both best efforts delivery and **end to end quality of service delivery**.

3.2…to ensure an adequate return on investment in high bandwidth infrastructures, operating agencies shall negotiate commercial agreements to achieve a sustainable system of fair compensation for telecommunications services and, where appropriate, respecting the **principle of sending party network pays.**

[3] see BoR (12) 33.

[4] ETNO Paper on Contribution to WCIT, ITRs Proposal to Address New Internet Ecosystem, ETNO Contribution (2012/09)

[5] BEREC Common Statement on Next Generation Networks Future Charging Mechanisms / Long-Term Termination Issues (BoR 10 24 rev1), ERG Common Statement on Regulatory Principles of IP-IC/NGN Core ERG (08) 26, Report on IP-Interconnection (ERG 07 09)

[6] See BoR (12) 33, which explains this value chain.