Internet Interconnections

Proposals For New Interconnection Model Comes Up Short

Key to the continued success and growth of the global communications infrastructure is to continue to allow the Internet to evolve in response to technological changes, innovation, and usage patterns.

Introduction

Recent contributions to the debate surrounding revision of the International Telecommunication Regulations (ITRs) have identified two network interconnection issues that, it is argued, are best addressed through new international regulation. These new international regulations would be implemented in the revised International Telecommunication Regulations (a binding Treaty), and implemented across legal and regulatory jurisdictions in ITU Member States (countries). These Internet interconnection issues are:

1. Whether to deploy an end-to-end quality of service delivery mechanism with pre-defined end-to-end performance objectives, and;
2. Whether to establish a system of settlements between network operators based on sending-party-network-pays.

This paper assesses whether there is a need to address either of these matters, which are currently handled through private commercial and technical agreements, in a global, intergovernmental treaty like the ITRs. It is our assessment that placing treaty-based interconnection obligations on infrastructure providers is not a constructive way to make progress. Further, we conclude that the full range of impacts that these proposals might produce has not been well documented. In our judgement, this proposed new interconnection model runs the serious risk of fragmenting the Internet.

Executive Summary

- **The Internet: a network-of-networks –** The Internet is composed of thousands of independently owned, managed and operated networks that interconnect with one another in a sparse mesh. End-to-end (E2E) traffic flow is enabled via a series of interconnecting customer, provider, transit and peering relationships between network operators, utilizing underlying telecommunications infrastructure and services.
- **The Internet interconnection market is diverse –** The market for Internet interconnectivity services between network operators is evolving and diverse (see BEREC). There are no technical barriers to Internet service providers (ISPs) making whatever sustainable interconnection arrangements they wish to make, provided they can get the other party or parties to the connection to agree to the terms.
- **Increasing costs for all ISPs** – Technical mechanisms required to support a generalised sending-party-network-pays settlement framework for the Internet in which one can measure both traffic volume and destination are extremely complex. Multiple new additional network components would be required to send, collect, and analyze the traffic data [ITU] and this settlement framework will, therefore, be extremely expensive to implement.

- **Sending-Party-Network-Pays jeopardizes competition –** Mandating a sending-party-network-pays system of settlements between and among network operators would put at risk the diversity and competition of the communications services marketplace that gave rise to the commercial Internet in the first place. This settlement regime would also jeopardize the distributed nature of the current Internet infrastructure by reducing choices over how and where traffic is sent.
- **Significant economic uncertainty –** There is no detailed, objective economic analysis of the sending-party-network-pays settlement proposal that indicates it will have a beneficial effect on the economic value of the Internet. Coupling sending-party-network-pays to the Internet will mean there is no way for content providers, transit providers or access ISPs to reliably predict what their settlement expenses will be.
- **Implementing end-to-end quality of service –** Currently, there are no technical barriers that prevent ISPs from deploying Quality of Service (QoS) enabled services *within* their networks. The barriers to deploying such services globally, between networks, are both technical and commercial. There is no need to introduce intergovernmental treaty text, that, if ratified, would lock in place a particular technical and commercial approach and thereby limit the future potential of the communications infrastructure. Key to the continued success and growth of the global communications infrastructure is to continue to allow the Internet to evolve in response to technological changes, innovation, and usage patterns.
- **Unintended consequences –** Contributions to the debate are insufficiently detailed to be absolutely clear about what the impact might be on end users if the proposals were implemented. Possible outcomes that have not been adequately explored include higher costs for end-users, selections by content and transit providers about where to send information which could impact information flows, additional incentives to misroute traffic for fraudulent purposes, fewer interconnection opportunities for developing countries, and fragmentation of the Internet. In short, these proposals have not considered the range of possible outcomes including the possibility of a much starker 'digital divide' than we already have today, potentially harming developing countries.

Given these factors, we do not believe that a new treaty-based global regulatory scheme that seeks to regulate how networks are managed, how networks are architected and configured, and how commercial agreements between network operators should be conducted is either necessary or appropriate. In fact, the introduction of treaty-based regulation in all of these areas may have significant and unpredictable negative consequences for the growth and evolution of networks.

Interconnectivity – Compensation and Commercial Arrangements

The Internet is composed of thousands of independently owned, managed and operated networks that interconnect with one another in a sparse mesh [PCH]. Interconnecting links come in one of two basic forms, either *transit* or *peering* (see **Figure 1**).

Transit is typically a bilateral agreement where an ISP provides full connectivity to the Internet for upstream and downstream transmission of traffic on behalf of another ISP or end- user including an obligation to carry traffic to third parties.

Peering agreements are where two network providers agree to exchange traffic between each other's customers (but not use each other's transit connections). Peering is the act of exchanging traffic with a peer [Tinka]. Peering provides network operators with a way to reduce the amount of traffic they have to send over potentially more expensive transit links. Networks normally implement peering when the cost of moving traffic via peering is less than moving traffic via transit.

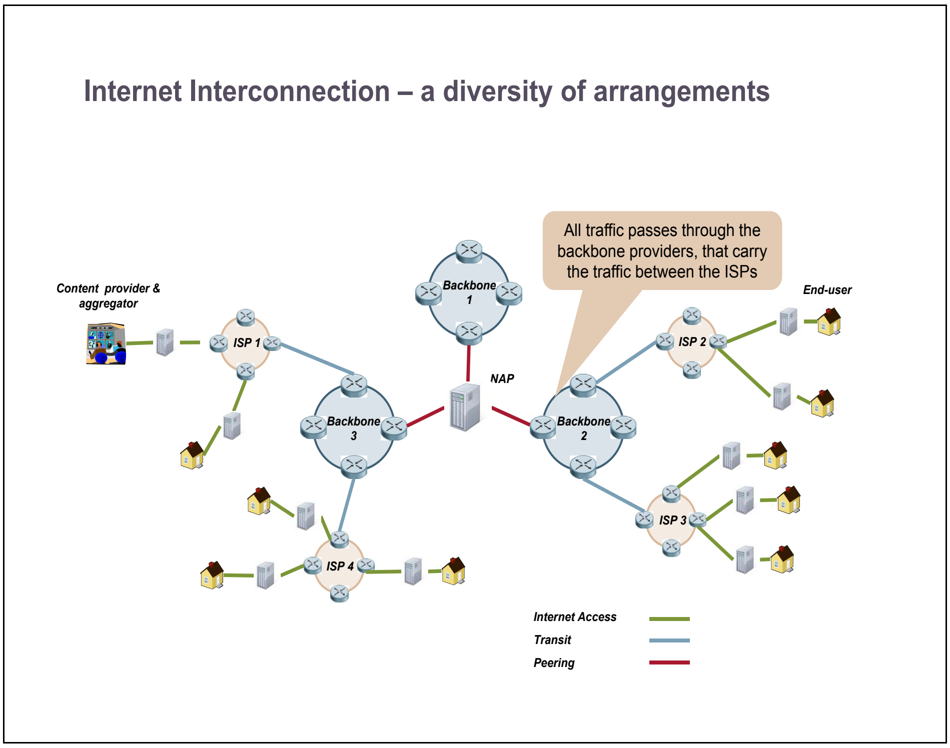

Figure 1: Internet interconnection (courtesy Analysys Mason): The Internet is a collection of networks and one interconnects arrangement is typically just one part of the complete end-to-end path between two end-points. By design there is normally at least one peering relationship in the path between two end-points.

While transit agreements are governed by normal customer-supplier commercial contracts, and may contain service guarantees, peering agreements are typically informal supplier-to-supplier agreements. In the vast majority of cases, peering arrangements are 'settlement free' – no money changes hands – because it is commercially mutually beneficial for both parties to interconnect. In some cases, asymmetric peering terms apply where one party compensates the other for carrying traffic either as part of a paid peering arrangement, or in order to meet minimum peering requirements [PCH]. Neither party to a peering agreement is restricted by regulation to the terms that they can request, nor are they obliged to accept any terms they find disagreeable.

Internet routing decisions are made on a per-packet basis and there are no fixed forward or reverse path network hierarchies to make charging outcomes predictable.

The commercial arrangements available to network operators today are therefore evolving and diverse[1] (for a richer view of what real-world Internet interconnection looks like, see [Faratin]). In simple terms, however, the following three arrangements are common:

- pay another operator for transit;
- peer on a settlement-free basis with other network operators, or;
- peer on a paid basis where minimum requirements for settlement free interconnection cannot be met.

Network operators typically employ a combination of peering and transit to engineer the most efficient and cost-effective solution for their needs. This evolving diversity of interconnection options and associated commercial arrangements provides operators with flexibility to pay for (and charge for) the connections they need to efficiently operate their networks at the commercial rates supported by a competitive market for Internet services.

As stated above, settlement-free peering arrangements are the most common type of peering arrangement, and are a well-established feature of the Internet ecosystem. After commissioning the most comprehensive survey of its kind ever undertaken, the OECD found that,

> "…the terms and conditions of the Internet interconnection model are so generally agreed upon that 99.5% of interconnection agreements are concluded without a written contract. That these "rules of the game" are so ubiquitous and serviceable indicates a degree of public unanimity that an external regulator would be hard- pressed to create. The parties to these agreements include not only Internet backbone, access, and content distribution networks, but also universities, NGOs, branches of government, individuals, businesses and enterprises of all sorts—a universality of the constituents of the Internet that extends far beyond the reach of any regulatory body's influence." [OECD]

The Internet architecture does not recognise national boundaries. Therefore, it is our view that an attempt to codify a new international model for commercial arrangements and agreements in a global treaty that imposes such boundaries is the wrong approach. This approach risks the flexibility and diversity that has evolved in the global Internet interconnection market to date, where traffic moves via the most commercially efficient route possible.

Sending-Party-Network-Pays

The proposed settlement model of sending-party-network-pays is imported from the telephone network-settlement model of 'calling-party-network-pays'. In the case of the telephone network, the 'sending party' is the person placing the call, or their network operator. In telephony, a 'call' initiates a reservation of network resources – the network reserves resources to support the subsequent communication between the initiator of the resource reservation and the called party. In other words, the requestor for the resource reservation is charged for the resources that have been reserved – the caller is charged for making a call. Calls are often provisioned through definable circuits and are visible to telephony networks due to expensive instrumentation of circuit switches. It is therefore feasible to interact with peer telephony operators using this notion of 'a call' as the currency of interaction. The Internet, however, was never designed to work on a circuit-based or circuit-switched system.

In a packet-based network such as the Internet, the network is simply unaware of the state of applications that are layered on top of it. There is no reservation of resources associated with an

---

[1] From the Body of European Regulators for Electronic Communications (BEREC) - http://erg.eu.int/doc/consult/bor_12_33_ip_ic_assessment.pdf. "The Internet ecosystem has managed to adapt IP interconnection arrangements to re- flect (inter alia) changes in technology, changes in (relative) market power of players, demand patterns and business models."

Internet communication in the way that there is for a telephone network call, no switched circuit, and no fixed path for packets to travel. Indeed, the Internet does not know whether you are making a voice call, downloading a video, or sending an email. The packets are sent on a "best effort" basis, which means there are no guarantees that packets will arrive within a certain time, or indeed, at all. Said another way, the difference between a telephone call and an Internet session might be seen as follows:

- **In a circuit-switched telephone call**, the circuit is opened and maintained as such throughout the call. So, even if both parties are saying nothing to each other, the circuit is open and the resources are reserved. This is why people have often been charged for phone calls on a per-minute basis: there's a cost for each minute that the circuit connection is established, even if the parties are not talking to each other.
- **In an Internet session**, however, there is no circuit. Packets only consume network resources for the time it takes to transmit them along the path to their destination. The forward and reverse paths taken by packets relating to a specific application session are not necessarily symmetric. Packets usually travel over three-to-five networks to reach their destination.

For this reason, we can see that on the Internet a 'call' is an abstract notion – it might be a property of the application (e.g., a session in a web request) or the application may never even have the notion of a call (e.g., transaction-based applications like the Domain Name System). Furthermore, since there is no established circuit, Internet routing decisions are made on a per-packet basis and there are no fixed forward or reverse path network hierarchies to make charging outcomes predictable.

The policy whereby a network hands traffic off to other networks at the earliest possible opportunity is called "hot potato routing" [Tinka]. In the example in
Figure 2, an end user connected via an ISP to Backbone 1 is requesting content from a content provider connected via another ISP to Backbone 2. The following steps occur:

1. Backbone 1 exchanges the request with Backbone 2 at the nearest IXP (Internet Exchange Point), which is IXP 1.
2. Backbone 2 then carries this traffic to the content provider.
3. The return traffic is then transferred back to Backbone 1 through the nearest IXP (IXP 2).
4. Finally, Backbone 1 carries the return traffic to the end user

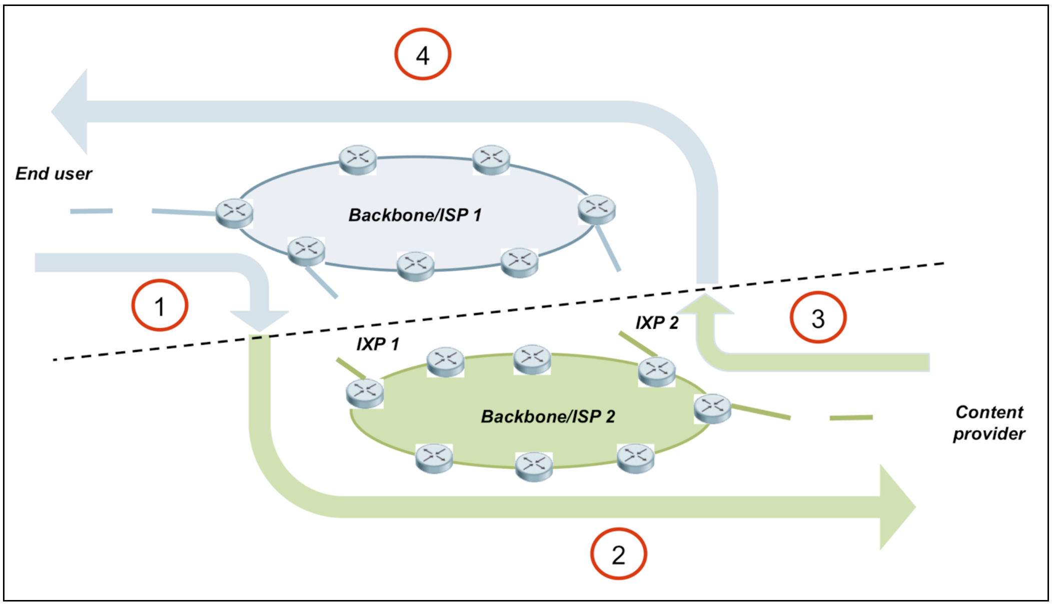

Figure 2: Hot Potato Routing (courtesy Analysys Mason) - Hot-potato routing is a scenario where a network operator wants to hand traffic off to other networks as quickly as possible. This is mostly typical of small networks that want to limit their costs by keeping as much traffic away from their network as possible. It is also an advantage for large networks that want to find the closest and quickest exit for traffic that will ultimately get off their network to reach its destination. The opposite is cold-potato routing [Tinka].

In a sending-party-network-pays context, there would be no way for content providers, transit providers or access ISPs to reliably predict what their settlement expenses will be. If a host on a network that imposes relatively high charges for receiving traffic requests Internet content from a remote network, then all network operators in the return paths may be charged.

Because the bilateral model of sending-party-network-pays or “sender pays” that is common in traditional telephony or mobile-settlement systems does not readily accommodate the Internet’s multi-party transit network system, it cannot be mapped to the Internet as we know it. Simply said, retro-fitting a “sender pays” settlement regime to the Internet is not possible without extensive changes to the infrastructure of the global Internet. In addition, the “sender pays” model could adversely impact the technical and commercial environment in developing economies that need to grow their networks. Ideally, countries would use mechanisms that work well in their local environments rather than globally mandated terms that aren’t well suited for the local context. Finally, as with any economic ecosystem, these interconnection arrangements have succeeded because the mutual interests among and between companies, Internet users, regulators, and civil society help to drive down costs for interconnection. Proposals by a few carriers to change this model may have the deleterious effect of altering the spirit and good-faith cooperation that has existed since the Internet’s inception. In short, these kinds of proposals to change this model would have far-reaching implications that need to be fully understood by *all* parties and weighed against the loss of the current spirit and good-faith cooperation that has existed since the Internet’s beginnings.

There are also significant public policy considerations if the regulatory system imposes a telephony-based settlement system on the Internet. For example, a sending-party-network-pays model could make it cost-prohibitive for content providers to send traffic to high-cost/low-revenue destinations. In theory, some content providers may determine that some destinations are simply too expensive to send any traffic to. Indeed, a requirement for content providers, wherever they are located, to predict the cost/benefit trade-off of deploying a new application or service in specific regions would have a chilling effect that could choke off the kinds of innovations that users expect in the modern communications age.

This might mean that the most remote destinations – and countries that are still developing their infrastructure – would be negatively affected because some content may simply be unavailable to them. Similarly, small content providers may find it very expensive to launch online content and to reach a global audience. Thus, a sending-party-network-pays model could lead to the international and regional fragmentation of the network into low-cost, content-rich regions and high-cost, content-poor regions.

Sending-party-network-pays could therefore reinforce and make much worse the existing ‘digital divide’, potentially locking-out some developing economies through prohibitive costs. In addition, these proposals do not seem to have weighed the possibility and consequences of importing the compensation schemes, scams, and arbitrage that plague the traditional telecommunications operating model to the Internet. We believe that the consequences could deepen economic challenges, increase the digital divide, and compound economic uncertainty. Finally, because an ITR provision governing Internet interconnection agreements could be implemented in myriad ways across the 193 ITU member nations, this proposed framework risks forcing a rising tide of inconsistent, unpredictable and potentially unworkable laws and regulations to a global Internet market that functions today precisely due to the fact that thousands of networks interconnect using informal agreements and open standards to facilitate that interconnection.

Sending-party-network-pays could therefore reinforce and make much worse the existing ‘digital divide’.

We expect that this would be a very negative outcome for the health of network development and the growth of the communications sector…

Content Delivery Networks

Recent studies [IXP] [UNESCO] have shown the nexus between economic growth and infrastructure and local content development. A recent trend in network architecture is to make content available locally, usually through some form of content replication and distribution system, such as a Content Delivery Network (CDN) (see **Figure 3**). This model has proven to be beneficial in reducing the need for content to be transmitted and re-transmitted over expensive long-haul links. CDNs can also improve the performance of applications due to reduced transmission latency.

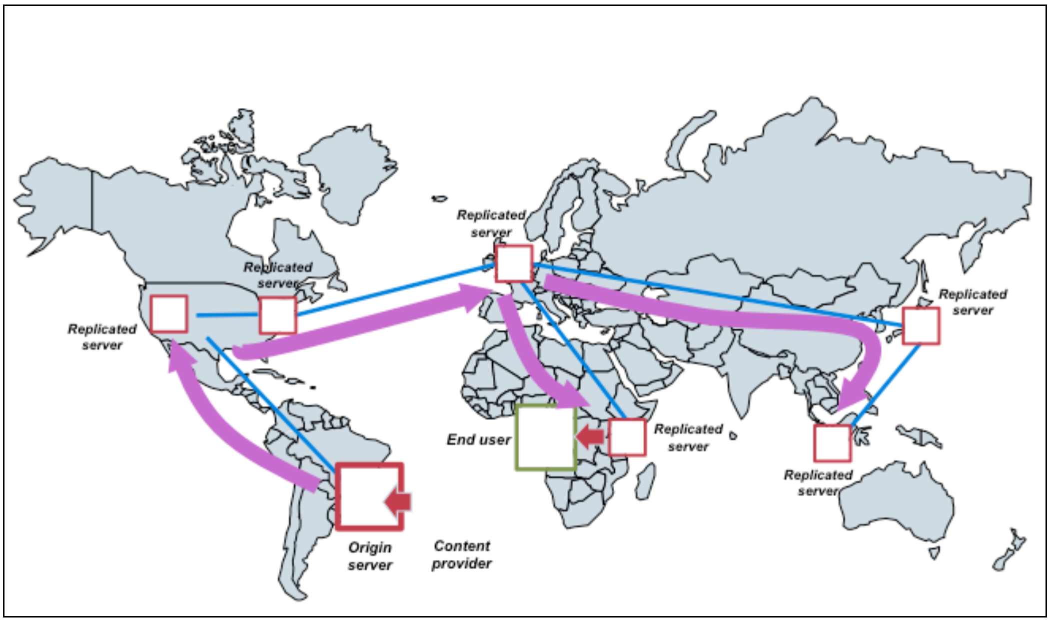

Figure 3: Content Delivery Network (courtesy Analysys Mason)

One particularly harmful consequence of the sending-party-network-pays settlement regime is that it may create financial disincentives for networks to invest in the network equipment required to cache content locally[2]. Under a sending-party-network-pays-model, the terminating ISP might maximise revenues from the inefficiently routed, but higher volume inbound traffic. As a consequence, content providers might be incentivised to filter high-cost remote destinations. This would lead to a vicious circle that stunts network development in developing economies where the need for investment in more connectivity and local content is the greatest.

If we imagine a world where a sending-party-network-pays settlement regime is in place, networks where Internet content servers were most prevalent would have to pay settlement fees to networks where such servers were less prevalent. In practical terms this would decrease the incentives to build local infrastructure thereby reducing the volume of inbound traffic and revenues associated with it.

As we’ve already identified, there are no rules in the Internet interconnection market today that prevent Internet operators from negotiating commercial agreements intended to ensure fair compensation for services rendered. By defining the terms of negotiation for Internet interconnection in the treaty itself, the sending-party-pays settlement regime would actually reduce the options available for negotiation, eliminate flexibility, and diminish diversity in the Internet interconnection market. We expect that this would be a very negative outcome for the health of network development and the growth of the communications sector, not to mention the technical challenges that would need to be overcome in order to implement such changes.

[2] Note: a cache is a dedicated network server that stores Internet content (e.g. Web pages or video content) retrieved by users, in order to serve future requests for the same data more quickly. The first time a user asks for a piece of content, such as a YouTube video, it is delivered from the international server to the user, while also stored in the local cache – subsequent users will then be served from the cache, saving on the cost and latency of accessing the international server.

Internet traffic is fundamentally different in character to circuit-switched telephony traffic and attempting to retrofit the accounting regime from the telephony world to the Internet is, in our view, inconsistent with the nature of the Internet's interconnectivity and the efficiencies that are inherent in the existing commercial arrangements. Thus, we find no merit in the proposals to incorporate a sending-party-network-pays settlement regime into the text of the ITRs.

Best Effort Internet Service

As mentioned above, the Internet relies on the concept of best-effort packet delivery. This implies no guarantees for considerations such as latency or packet loss. This does not mean that an individual network that is part of the Internet is not managed and is not engineered to meet performance objectives. As a network of networks, the individual networks that comprise the Internet agree to do their best to exchange, route and deliver packets presented to them given the available resources. When resources are congested, packets get dropped, and end hosts treat dropped packets as a signal to send less data, thereby preventing ‘congestion collapse’ of the network. For most applications, dropping packets is not a problem because they can simply be re-sent with little or no consequence. For example, if a user sends an email that has “lost” packets on the way, the systems typically notify each other and “self-heal,” meaning that the packets are re-sent. Concretely, this might mean that an email takes longer to be delivered than before (because of the need to resend some packets), but in practice this happens automatically without any action from the user.

Although consumers are tolerant of lower quality in exchange for lower costs, it has been argued that best-effort packet delivery is insufficient for inelastic applications that have a low tolerance for packet loss and high or variable latency. Examples of such applications include Internet telephony, video-conferencing, and audio and video streaming. Over time however, we have seen significant advances in network quality as bandwidth has increased and new “broadband” access network technologies have been deployed. We have also learned that many such inelastic applications in fact work tolerably well over a best-effort network [RFC5290].

At the same time, a number of these applications or services have been provided over QoS-enabled, private, managed IP networks that use the same IP-based technology that underlies the Internet, but that do not interconnect with the Internet. Otherwise known as premium Internet services, these have been expensive in both implementation and management. The lack of guaranteed end-to-end QoS on the Internet can be viewed as a positive feature – diverse applications can be supported without the additional expense and management complexity.

Also, network management options are becoming increasingly sophisticated. There are no technical obstacles that prevent a network operator using Quality-of-Service (QoS) mechanisms to implement different classes of service *within* their network today[3]. Indeed, this technology has been standardised, implemented and deployed in closed networking environments for many years. The challenge in replicating this experience over the Internet comes with trying to deploy QoS mechanisms in an inter-provider topology where there are many different players, as depicted in **Figure 4** and **Figure 5** below. [Huston] [QoSFiction].

[3] The Internet Society's principles with respect to network management and open internetworking may be found here: http://www.internetsociety.org/open-inter-networking-getting-fundamentals-right-access-choice-and-transparency

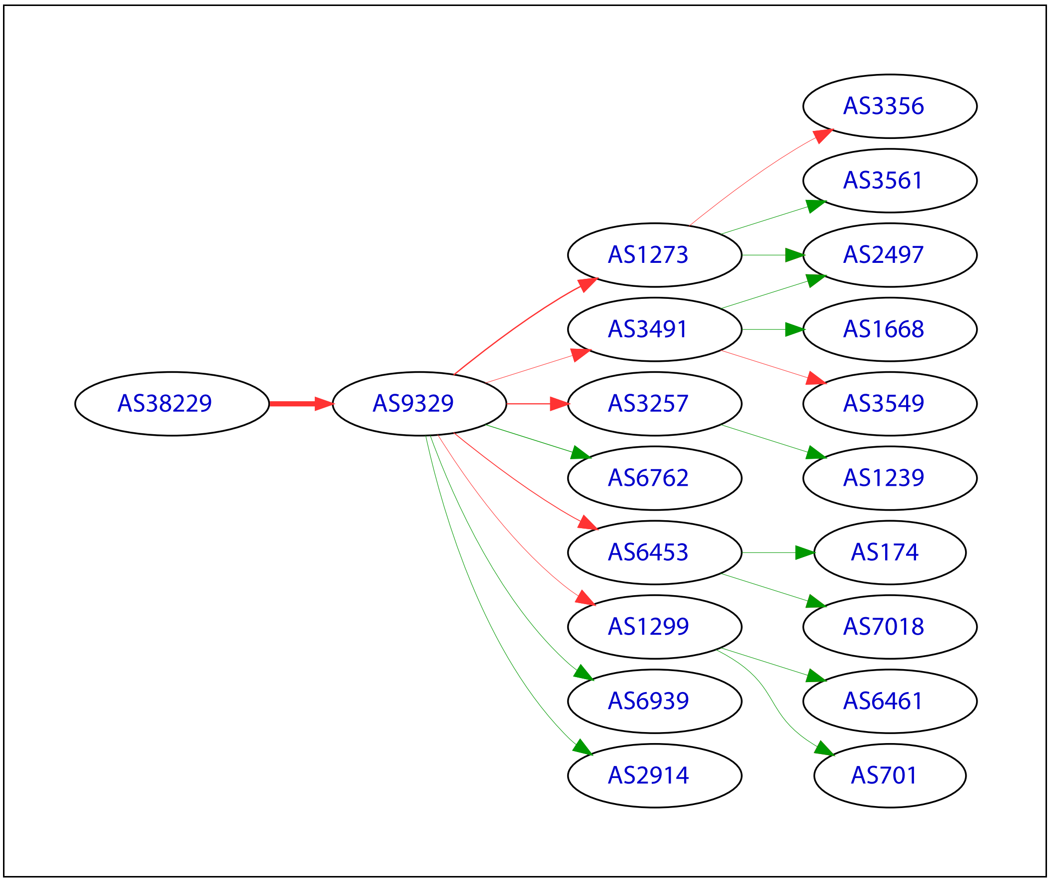

Figure 4: **Interconnection complexity** (courtesy Hurricane Electric): In practice, interconnection arrangements can be extremely complex. The first example shows IP relationship (transit or peering) Sri Lanka Education & Research Network, NREN as of 19 July 2012 (http://bgp.he.net/AS38229). Red lines indicate transit relationships; green lines indicate peering relationships.

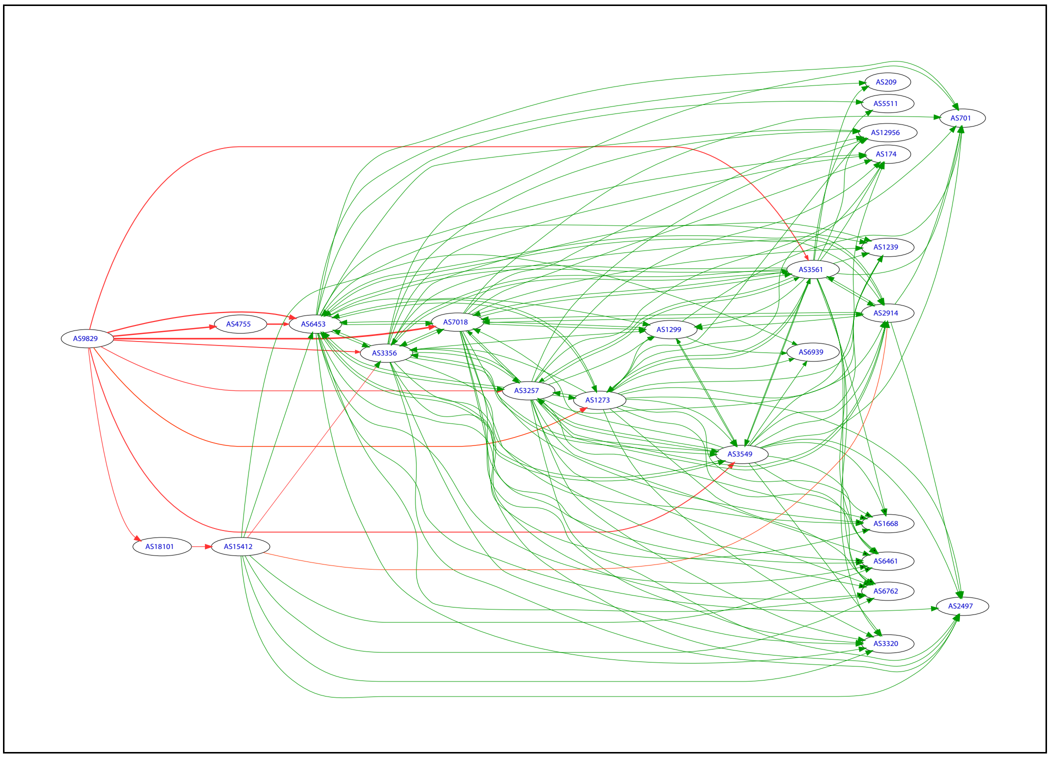

Figure 5: Interconnection complexity (courtesy Hurricane Electric): In practice, interconnection arrangements can be extremely complex. This example shows IP relationships (transit or peering) for Bharat Sanchar Nigam Ltd from India as of 19 July 2012 (http://bgp.he.net/AS9829). Red lines indicate transit relationships; green lines indicate peering relationships.

The key is to continue to allow the Internet to evolve in response to technological change, innovation, and changing user patterns.

The ITRs are the wrong place to attempt to re-architect and re-define Internet interconnection.

Delivering a 'premium' service based on pre-defined end-to-end performance objectives would require multilateral provider assurances to deliver end-user benefits. Such assurances are not a general feature of the Internet today due to the difficulty in achieving agreement on QoS parameters between such a large number of providers and the enormous complexity of the technical mechanisms involved. While tremendous adaptability, flexibility, and resiliency are its core, the Internet evolves in phases as individual operators make autonomous business decisions (see **Figure 5** above).

The key is to continue to allow the Internet to evolve in response to technological change, innovation, and changing user patterns. If consumers demand QoS, the diverse actors that build and operate the Internet have proven that they can respond without the need for global government mandates or regulation.

Given the current complexities and technical challenges underlying the delivery of QoS over the Internet, it seems to us that imposing premium service obligations on all network operators through international regulation carries the serious risk of introducing tremendous cost inefficiencies into the communications sector without a clearly identified benefit for Internet users. That is why it is critical that these changes occur in a manner that allows technical and network management decisions to be made by the entities operating those networks, and in response to user demand in a commercial environment, and not through unprecedented global government intervention. From this perspective, introducing text into the ITRs relating to an end-to-end quality of service delivery mechanism is highly undesirable and could have serious unintended consequences for developing economy operators that are trying to deploy networks at lower costs.

Conclusion

The Internet is made up of over 40,000 networks exchanging over 425,000 globally unique routes[4], with 2.2 billion users and 500,000 new users each and every day[5]. With over 600 Million web sites[6] and over $600 billion US dollars in annual ecommerce[7], we live in a world where you can find unprecedented amounts of information with ease. In just 20 years, the Internet has grown to be critical to both the economic output of many countries and the daily lives of billions of people, and note that all of this astounding growth happened with the governance of the Internet in the hands of those who own it, manage it, develop it and use it, via a host of multi-stakeholder groups.

We see no evidence or sufficient analysis that demonstrates that mandating the deployment of an end-to-end quality of service delivery mechanism would result in any net benefits for consumers or the economy. Technical mechanisms exist that allow operators that are able to make a business case for QoS deployment to do so.

An inter-governmental treaty such as the ITRs is therefore the wrong place to codify or lock-in business models or technologies, particularly those that run the risk of fragmenting networks or penalizing countries that request content that some networks do not want to deliver. The ITRs are simply the wrong place to attempt to re-architect and re-define Internet interconnection. Instead, policy makers, working with their Internet stakeholders, should focus on developing the appropriate policy frameworks that foster competition and support the continued growth and evolution of the Internet in their countries. This allows countries to set policies that reflect local market conditions rather than locking in a one-size-fits-all, global regulatory approach that may have broad, unintended consequences.

---

[4] potaroo.net
[5] internetworldstats.com
[6] netcraft.com
[7] JP Morgan, 2009

References

[ITU] Recommendation ITU-T D.50 – Supplement 1: General considerations for traffic measurement and options for international internet connectivity

PCH] Woodcock, B., and V. Adhikari, “Survey of Characteristics of Internet Carrier Interconnection Agreements”, Packet Clearing House, May 2011

[Tinka] Tinka, M. “Peering Jargon” (http://www.internetsociety.org/terminology-bi-lateral-peering-mark-tinka)

[IXP] “Assessment of the impact of Internet Exchange Points – empirical study of Kenya and Nigeria”, Internet Society, 2012 (http://www.internetsociety.org/ixpimpact)

[UNESCO] “The Relationship between Local Content, Internet Development and Access Prices”, Internet Society, OECD, UNESCO, 2011 (http://www.internetsociety.org/localcontent)

[RFC5290] Floyd, S., and M. Allman, "Comments on the Usefulness of Simple Best-Effort Traffic", RFC 5290, July 2008 (**http://www.ietf.org/rfc/rfc5290.txt**)

[QoSFiction] Huston, G., “Quality of Service – Fact or Fiction?”, Internet Protocol Journal, Vol.3 Number 1, March 2000 (**http://www.cisco.com/web/about/ac123/ac147/ac174/ac197/about_cisco_ipj_archive_article09186a00800c8314.html**)

[OECD] “Internet Traffic Exchange: Market Developments and Policy Challenges”, OECD Working Party on Communication Infrastructures and Services Policy, October 2011

[Huston] Huston, G., “The QoS Emperor's Wardrobe”, June 2012 (**http://www.potaroo.net/ispcol/2012-06/noqos.html**)

[Faratin] Faratin, P., D. Clark, S. Bauer, W. Lehr, and P. Gilmore, "The Growing Complexity of Internet Interconnection", Communications & Strategies, number 72, 4th quarter, 2008

Internet Society
1775 Wiehle Avenue,
Suite 201
Reston, VA
20190-5108 U.S.A
Tel: +1-703-439-2120
Fax: +1-703-326-9881

Galerie Jean-Malbuisson 15
CH-1204 Geneva
Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445

Email: info@isoc.org
http://www.internetsociety.org

About the Internet Society

The Internet Society (ISOC) is a non-profit organization founded in 1992 to provide leadership in Internet related standards, education, and policy. With offices in Washington, D.C., Geneva, Switzerland, and 10 other countries, it is dedicated to ensuring the open development, evolution, and use of the Internet for the benefit of people throughout the world. ISOC is also the organizational home of the Internet Engineering Task Force (IETF), the Internet’s premier technical standards body.