

September 2022  
Geoff Huston

## Sender Pays

In September 2012 ETNO, the European Telecommunications Networks Operators' Association, or most notably Deutsche Telekom and France Telecom and fellow legacy telcos in Europe published a contribution to the 2012 World Conference in International Telecommunications (WCIT-12) with a proposal for regulatory reform that in ETNO's words would compel content providers to directly contribute to the costs on Internet communications infrastructure. Or, in their words, "advocating for an adequate return on investment based, where appropriate, on the principle of sending party network pays" (<https://etno.eu/datas/itu-matters/etno-ip-interconnection.pdf>).

The proposal excited some reaction at the time, with varying levels of hyperbole. There was the Centre for Democracy and Technology's dire warning that the "ETNO proposal threatens to impair access to [an] open global Internet" ([https://www.cdt.org/wp-content/uploads/pdfs/CDT\\_Analysis\\_ETNO\\_Proposal.pdf](https://www.cdt.org/wp-content/uploads/pdfs/CDT_Analysis_ETNO_Proposal.pdf)) which was pretty representative of the reaction from much of the Internet sector.

At the time there were no changes arising out of WCIT-12, and the matter would've remained as a simmering undercurrent in the industry, were it not for Korea.

### **Korean Content Wars**

Back in 2012 so-called "smart TVs" were the new thing, and Samsung launched a high-definition video streaming service in South Korea. The response from Korea Telecom (KT) was a rapid escalation, by blocking these units from accessing KT's broadband access network. Considering that KT was the country's major broadband access provider and Samsung was (and still is) a major retail brand this was always going to excite a public reaction. Samsung's response was to invoke the principles of Network Neutrality, whereby, they claimed, consumers should be able to use network services without any active discrimination by the network provider, and also casting doubt on the argument that the streaming service was an excessive consumer of network capacity in any case. The Korea Communications Commission stepped in and called KT's actions "inappropriate". KT backed off with the network blocks on these televisions and the inexorable rise of high-definition content streams continued in Korea, and all over the Internet, in the ensuing decade.

But the issue never went away, and it had a reprise in 2021 when SK Broadband (SKB) launched a legal action, claiming that Netflix should pay the costs of supporting increased traffic loads in response to a surge of SKB's customers streaming Netflix content. The Seoul Central Court ruled that SKB has legitimate grounds for compensation and the amount is a matter of negotiation between SKB and Netflix. Some Korean lawmakers have spoken out against content providers who do not pay

for network usage despite generating large traffic volumes. Clouding the issue in Korea was the additional factor that some local content providers were already paying Korean ISPs for content stream access, so this dispute was easily cast as a battle between local enterprises and an overbearing US giant refusing to do what the local players are already doing.

So, the entire set of issues of network neutrality, interconnection and settlements, termination monopolies, cost allocation and infrastructure investment economics are back with us again. This time it's not under the banner of "Network Neutrality," but under a more directly confronting title of "Sender Pays". The principle is much the same: network providers want to charge both their customers and the content providers to carry content to users.

## Sender Pays

Let's pause for a second to look at the legacy of "Sender Pays" in the context of this debate.

When a service is constructed using diverse components, then the way in which service revenues are distributed to the various suppliers of the components of the service can follow a number of quite distinct models. There is the incremental payment model where the customer directly pays each service provider for their service independently of the other services. (Of course, imagine the privileged position of the final provider the customer has already invested heavily in the provision of the service and the final provider is in a position to charge a far higher price as the transaction has already accumulated a sunk cost. Then there are various forms of revenue redistribution models where the revenue per transaction is distributed by the service orchestrator to the various suppliers according to their inputs to support each transaction. The service orchestrator can bill the initiator of the service transaction or bill the service host. And there are various "bill and keep" models where a set of service providers use each other's services by essentially donating their service to each other. This allows each service provider to bill their customers for the entire service and rely on a "knock for knock" model where the foregone revenue for the donated services matches the opportunity revenue gained in the efficiencies not having to redistribute the fees paid by the customer for the compound service.

If we start with the mail service, the initial models were based on "receiver pays" where the message accumulates a debt on its progress towards delivery, and the debt is paid upon delivery of the letter. The revenue could be redistributed to the other couriers following payment, but a simpler scheme is for each successive courier to purchase the letter from the previous courier. In this way there is a strong incentive to complete the delivery as the terminating courier may have already funded the letter's progress through the courier system.

The Royal Mail took two revolutionary steps in the nineteenth century that would make it the leader in postal innovation at the time. In 1840, Roland Hill, an enthusiastic campaigner for postal reform, led the Royal Mail's introduction of the Penny Post, a cheap single rate for domestic delivery of letters, replacing a complex distance and letter volume tariff. This step made the mail service accessible to a much larger population. At the same time, it was intended to result in a large-scale uptake of the use of private correspondence, allowing the mail delivery service to realise economics of scale and process. The second revolutionary step was a switch from the recipient pays to a sender pays tariff model. This step removed a troubling imposition on the mail service, caused by mail refusal. Under the recipient-funded service model, the refusal of a mail item meant that the postal service incurred the cost of delivery, as well as the cost of return, assuming that the refused mail was ever returned. These reforms had a profound impact. They generated huge volumes of postal traffic, facilitated the development of commerce and underpinned significant changes in social communication. This measure coincided with the Industrial Revolution, when people were moving from rural areas to the industrial cities and from one country to another, so these postal reforms eased to, some extent, the social pressures that arose from this large-scale transformation in nineteenth century society.

The telephone system followed this postal template. A telephone subscriber paid the full cost of a telephone call to the local telephone company. If the call involved another telephone operator to complete the call, then the calling operator paid the terminating provider an agreed rate for their part of the service. The inference from this model is that there is a transactional accounting model for the service, where the user is charged for each transaction, and it is this revenue that is used within the inter-provider settlement regime.

At one point in the telephone saga one large telco acknowledged that it cost the company some 21c to bill their customers 25c per local phone call. The shift to flat fees without usage tariffs is not only attractive to many customers, but also attractive to many operators as a means of stripping out huge costs from their billing systems as well as eliminating one of the major points of friction with their long-suffering customer base.

This also leads to the observation that a stable inter-provider settlement regime should be based on a uniform service tariff. How a service is billed to a user should determine how the service revenue is distributed between the various service providers who contributed to deliver the service. When they differ then there is an obvious opportunity for arbitrage and distortion.

While the sender pays inter-provider settlement regime was used for many decades, that does not mean that it was free from various forms of abuse. A national provider could raise offshore revenue by raising the local international call tariffs for outbound calls, making it attractive for international callers to call into the provider's network. This way the provider could use the inter-provider payments associated with call termination of these calls to raise revenue.

This opportunity was not lost on many national telcos and the resultant arbitrage of international call accounting settlement rates was often one that works to the financial disadvantage of AT&T in the US. It was no coincidence that when the Internet emerged AT&T lobbied US politicians and the administration to keep the Internet out of the clutches of the ITU-T.

AT&T clearly believed that the root of this problem lay in the regulatory framework that was promoted by the ITU-T. They believed that the one-country-one vote system biased the entire process in favour of smaller nation states and the outcomes were rarely ones that coincided with AT&T's interests. If the Internet was to be constructed on a platform of built through private sector investment, then the international governance structures should reflect the same primacy of private sector interests.

From the perspective of a telephone operator everyone is a customer.

## Enter the Internet

The Internet was originally conceived as a model that was, at least in terms of functionality, akin to the telephone network, but their differences are far more important than their similarities.

There is no "call" in the Internet. There is no confirmed transaction that initiates a network state that has attributes of distance and duration. Packets are far more casual than that. Packets may be delivered

to their intended destination, or not. They may have been solicited, or not. They may be intercepted by intermediaries or agents, or not. They may trigger other packets, or not. So, in the absence of a clearly defined transaction model of interaction between the Internet and its connected hosts, then a transactional accounting method to recover costs through transactions was never going to work. So, how was the Internet supposed to be funded?

Contrary to the utopian mantra of the late 80's the Internet is certainly not free, and ultimately there is only one source of revenue: its users like you and me. We each pay our local access providers some form of service tariff, typically in the form of a "whole of network" access tariff. We do not expect to be charged more if the packets are to be delivered to a destination on the other side of the world as compared to next door. We don't expect to be charged incrementally to receive packets, nor to send packets. The tariffs are based on an "access" model, where the service provider is providing a comprehensive access model. The retail differentiator is typically the access bandwidth, and in certain markets, notably mobile services, some count of aggregate traffic volume.

In the absence of a transactional model as a common retail structure then how do various providers financially balance their respective contributions to providing a seamless end-to-end experience to end users? The model used by the Internet is extremely crude, but quite effective.

When two networks directly interconnect and exchange traffic then either one network is a customer of the other and pays for all traffic sent and received, or the two networks regard each other as "peers" in the true sense of the word, and exchange traffic without any payment between them. From an individual network's perspective, the network pays for a "transit" service provided by its contracted upstream providers, it receives revenue from its customer networks, and it "peers" with other networks where it exchanges access to its customer base with the peer network's customer base without payment, but does not provide access to its transit providers, nor to any other peer networks. (Figure 1).

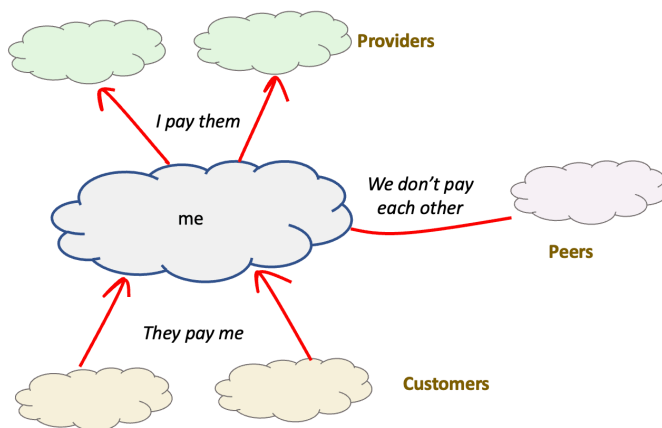


Figure 1 – Inter-provider settlements

Who determines who is the customer and who is the provider? Who determines which networks are peers of a given network? The Internet's answer is "nobody". The determination of roles is left to the networks themselves and the peering and interconnection regime functions as a market.

The result is that all viable paths between endpoints on the Internet have just two parts: a part where the packet's sender is funding the packet's transit, either directly or indirectly via customer/provider inter-provider relationships, and a second part where the packet's receiver assumes funding responsibility for the packet's transit (Figure 2)

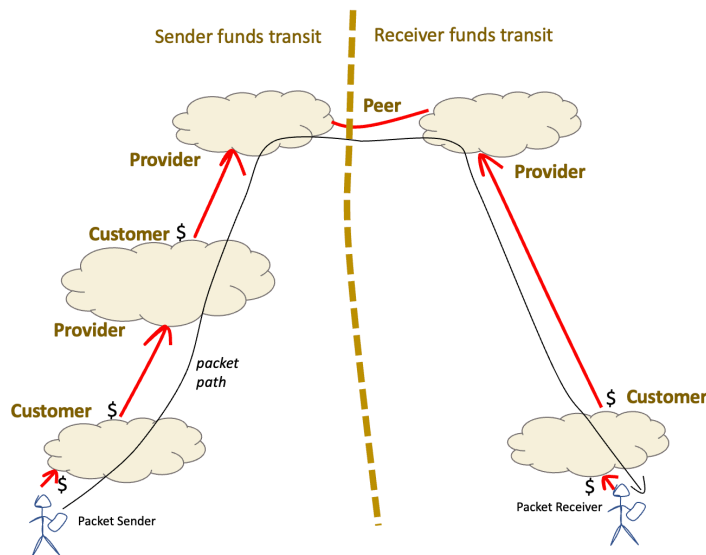


Figure 2 – Inter-provider path settlements

This Internet path property is sometimes referred to as “valley-free” paths. If a customer/provider relationship is an “upward” relationship, a provider/customer relationship is a “downward” relationship and a peer relationship is a “horizontal” relationship, then any viable path starts with a sequence of customer/provider hops (*up*), then at most one peer hop (*across*) then a sequence of provider/customer hops (*down*). All financially viable inter-provider paths can be presented as a “mountain” in this taxonomy. A “valley” would be where a packet is received by a network as a customer but then passed to a different transit network also as a customer (Figure 3). The valley segment is unfunded either by the sender nor the receiver.




Figure 3 – Inter-provider “valley” path

So, we now have a lexicon of *customer/provider* and *peer*, and an associated understanding of the flow of money across these two relationships. What about *paid peering*? This is a variant of a *peer* relationship where the *paid peer* receives the same access as a conventional *peer* but pays the provider for this access. It can be seen as a subset of a conventional customer relationship without access to the providers transit providers nor to any other *peer*, presumably at some discount to the normal customer tariff, so in many ways *paid peering* is just another instance of a *customer/provider* relationship, but without the stigma of being labelled as a customer!

## Service Provision

The internet interconnection model assumes that the interior of each of these networks are a collection of carriage circuits and packet switches. However, an ISP performs a number of other functions which can be thought of as part of the ISP's service, and these services are bundled into the retail product.

A good example of this implicit bundling can be found in the provision of Domain Name System (DNS) name resolver servers. These services are provided by the ISP and made available to the ISP's customers. They are typically not made available to the ISP's transit providers, nor to the ISP's peers. The cost of operating such services is met by the ISP's customer base through the ISP's tariff structure.

Such ISP-provided services are not limited to DNS services. Prior to the ubiquity of Google's Gmail service many ISPs offered mail hosting, and perhaps web hosting and related services. In some cases, these were bundled into the retail tariff, and in other cases were optionally selected services available at an additional tariff.

The next step in the evolution of the provider service model is the third-party provision of such bundled services. An ISP may pay a third party to provide DNS services, for example. Conventionally, this is seen as a service provided to the ISP, at a cost to the ISP, which is bundled into the ISP's retail offering.

However, the distinction between a third party acting as an agent of the ISP and a third party acting independently start to get blurred when the service model concerns third party content hosting and similar. Having content delivered directly to the ISP's customer base allows for a cheaper and faster service for both the ISP's customers and the content provider's content customers. The question naturally arises that if hosting a third-party content service within the ISP's infrastructure saves costs for the content provider, then rather than the ISP paying for the service and folding that into the ISP's bundled tariff, shouldn't the content provider be paying the ISP for access services in the same vein as any other customer? What makes the content service a special case here that it gains free access to the ISP's network? Why should a content provider be freeloading over the ISP network when all others pay for similar access?

There is no objective and clear answer to this question. Instead, the answer is the outcome of a negotiation between the parties and the outcomes vary based on the relative pressure and leverage each party can bring to the negotiation.

## Choke Points and Natural Monopolies

A "natural monopoly" is a type of monopoly that arises where the unique position of the incumbent is such that there are significant barriers to entry for any potential competitor. A sea port might be located in a particularly favourable point that makes it uniquely useful for trade, such as the location of Rotterdam at the mouth of the Rhine, or Copenhagen at the point where the Baltic Sea meets the North Sea.

There are various choke points in the Internet infrastructure which become natural monopolies. Exclusive use radio spectrum allocations for mobile service providers are a good case in point where the usable spectrum is a limited resource, and the spectrum auction process typically has just 3 or 4 operators each with exclusively licensed spectrum while other intending mobile operators have to negotiate arrangements to share the mobile infrastructure. The broadband access networks also can be considered as a natural monopoly given the high cost of competitive deployment of multiple access networks. Even a single broadband deployment seeks a household connection rate of more than 60% of households passed to make the deployment cost efficient. Multiple parallel deployments of broadband access infrastructure become extremely prohibitively inefficient very quickly.

How does this relate to the concept of "sender pays" in relation to the negotiations between access providers and content networks? To rephrase this in the terminology of Internet interconnections some access ISPs would like to position the content streamers as a "paid peer" and have them pay some termination fee to have their content reach the ISP's customers. The content streamers are asserting that



they have carried content across oceans and across continents and have delivered the content to the front door of the ISP, all at their expense, and it seems like a crude form of extortion for the ISP to then charge the content stream to carry the content up to the third floor!

It is a point of leverage for the access provider. For example, in Germany the three largest access providers serve some two thirds of the national user base, and there is no feasible way for an alternative mechanism to access these users, and no real prospect of churning these customers to a different provider. So, if the position from these three providers is one that demands some form of access co-payment from the content provider as a termination fee, then the content provider is likely to accede to these demands and pay the impost. To put it simply, there is no getting around a termination monopoly.

### **Negotiation or Blackmail?**

What if the content provider cannot reach an acceptable arrangement to directly host content servers within the access provider's network?

This is not an abstract question. If the access ISP takes the position that these third-party content providers should be customers of the ISP, then the content provider has little choice but to either enter into a customer arrangement or take a step back and offer the service via an external peer or transit network. The content streamer then is not paying the access ISP a termination fee and the ISP's customers can still access the streamed content. Problem solved. Right?

But that's by no means the end of the story. Part of the negotiation playbook as seen the access ISP deliberately allow this external peer or transit connection to reach the content provider to congest such that the service quality of the streamed content degrades significantly. This has happened in many contexts over the years, such as in the early 2000's when one large economy decided to force transit ISPs to pay a termination fee to access their national internet infrastructure. The forcing function was to only allow access only via under-provisioned access path unless the transit ISP was willing to pay the termination fee. The same technique is used today when a large access ISP deliberately pushes off-network content services to enter the ISP's network via an under-provisioned path which significantly degrades the customer experience when accessing the service.

This deliberate differentiated treatment of traffic from certain content sources in order to exert leverage over them falls under the general topic space of "Network Neutrality." In theory the customer base has contracted the Access ISP to deliver all content to the customer on equal terms. Under network neutrality framework the network should not treat traffic from different sources in different ways. This has been a high-profile issue in the United States, where the issue of enforceable measures to require public Internet networks to carry all traffic in a consistent and neutral manner hinges on whether an Internet service is defined as an "Information Service" as defined by Title I of the US Telecommunications ACT or as a common carrier under Title II of the same act. If ISPs are Title II common carriers, then the FCC has the power to enforce measures on these providers including behaviours consistent with the principles of network neutrality. The 5-member Commission of the FCC changes with each new Presidential administration, and the attitude of the FCC has reflected the seesaw of US political leanings in recent years.

Under FCC chair Tom Wheeler, the FCC voted in the 2015 Open Internet Order, categorizing ISPs as Title II common carriers, and making them subject to net neutrality principles. Under the Trump administration the FCC chair Ajit Pai, had the FCC reverse the 2015 Open Internet Order in December 2017, reverting ISPs to Title I information services. As part of an executive order issued in July 2021, President Biden has called for the FCC under the current chair, Jessica Rosenworcel, to undo some of the Trump-era changes around this Title I classification.

### **Negotiation Tactics vs Reality**

The conventional rationale from the access provider is that the demanded termination fees would allow the provider to cover the costs of incremental upgrades to the internal infrastructure of the access

network. Is this really the case? Or is this a case of opportunistic rent seeking on the part of the access provider?

A typical access ISP infrastructure is shown in Figure 4. The internal design of these networks is a collection of edge fan-outs, where each edge concentrator is connected inward to a Point of Presence (POP), which are themselves interconnected over a backbone set of trunk connections. Typically external connections to exchange points, transit upstream providers and peer networks are also made from backbone sites, so the edge network is dedicated for customer connections.




Figure 4 – Generic Access ISP design

Over the past decade streaming has taken over much of the traffic profile for a retail ISP. All CDNs offer settlement free peering and provide private interconnects over IXs irrespective of the size of the ISP (Figure 5a). Larger ISPs may be offered central cache servers or server clusters without cost to the ISP (Figure 5b). Large ISP may use CDN cache clusters, or use caches deployed closer to the edge of the internal infrastructure (Figure 5c). In this edge deployment scenario, it's unclear as to the real extent of the impact of this streaming content in the ISP's internal infrastructure, as the content streams are passed on demand only across the access tails rather than across the entire ISP's internal infrastructure.

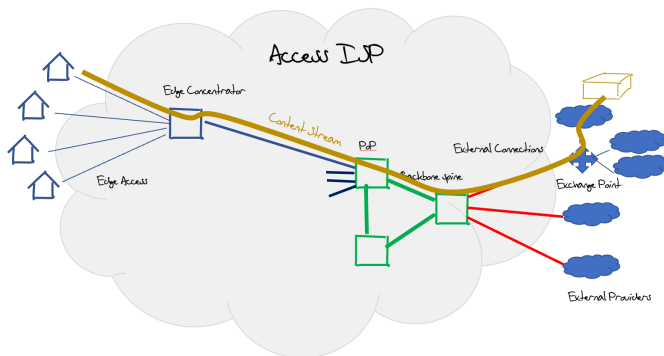


Figure 5a – External Content Streaming

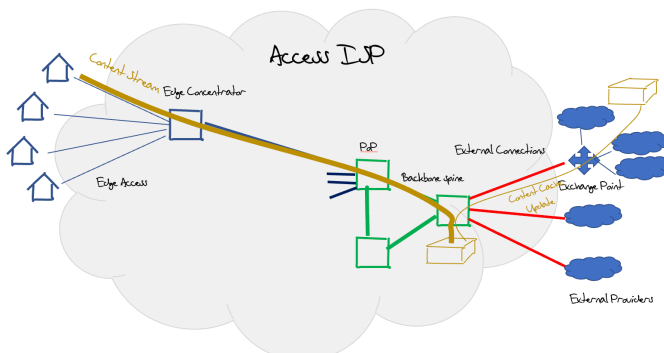


Figure 5b – Internal Content Streaming






Figure 5c – Edge Content Streaming

What is also unclear is the extent to which existing customer payments have ‘covered’ the ISP’s costs. Given that the major case to deploy a high-speed broadband network is to stream content, then it does seem to be either negligent of the access ISP sector if they failed to factor in this use pattern while at the same time deploying ever faster 5G mobile services and high-speed fibre broadband infrastructure. Or, more likely, this use case scenario was already factored into the access product from the outset and the claims that the streaming use by customers have created an unanticipated load on their network infrastructure to be specious at best, or, more likely, an outright lie.

The critical factor for residential broadband deployment is not the costs to pass a house (in USD that cost still appears to be some \$1,000 per household as part of a larger deployment), but the take up rate of the service. The cost of the tie line to the house is some \$800. If the take up rate is 50% then cost of the project rises by 55% as compared to a 100% take up rate. The costs associated with the take up rate of broadband dominate the economics of broadband deployment, and the marginal cost of bandwidth and volume are relatively minor cost factors. Accordingly, it’s challenging to substantiate the claims of significant costs on the part of the Access ISP to meeting the traffic demands of carrying streaming services.

## Consequences

The issue here is that in a relatively unconstrained negotiation the outcomes tend to favour the larger player. If the content provider is a major player with content that customers are demanding, then the ISP is under some pressure to reach an acceptable outcome or risk customer churn to a rival access provider. On the other hand, a major access provider is able to present an effective termination monopoly to the content provider and can demand a termination fee from the content provider as there are few alternatives to the content provider other than withdrawing from that market completely while the content streaming business is subject to a wave of competitive pressure from the various new entrants.

What improves the negotiating position for both parties is relative size. The greater the customer share the greater the leverage the access ISP can bring to the negotiation. The same applies to the content provider, and the larger the content provider in terms of market share the better its ability to capture a significant committed customer base, and then use this committed customer base as leverage in the negotiation with the access ISP.

The common outcome for both parties is increased pressure to aggregate and increase their size, or, in other words, continue along the inexorable path to centralization and large-scale industry consolidation, which is already a prominent issue for today’s digital world.

## Options for the Regulator

How should a regulator respond to this situation?

It’s clear that a detailed regulatory response that sets forth strict guidelines on how termination fees can be levied, and under what circumstances, stands a strong risk being overtaken by events in relatively short

order. Overly prescriptive directions have the risk of rapid obsolescence as the players, the services, and the technology options change.

The larger danger lies in potentially compromising the ability of the national economy to sustain continued private sector investment in broadband infrastructure through increased regulatory risk. There are very few economies where the construction of a broadband access infrastructure has been undertaken as a public sector activity, and if it's a private sector activity then each national environment is competing for investment funds with other national environments as well as other activity sectors. The capital returns on infrastructure investment in this sector have been generally rather poor, and this is reflected in the protracted exercises that make up various national broadband rollout programs.

However, sustaining a favourable environment for private sector investment in infrastructure is not the only objective here. An overly predatory private sector that exploits a monopoly position to the detriment of local consumers and local enterprises is equally a major liability for any economy. In this post-industrial social landscape, many economies are pinning their hopes for a sustainable future on a vibrant and healthy digital economy. A fundamental fracture in the relationship between content and carriage services imperils any such aspirations for a healthy and efficient digital economy at a national level.

The currently fashionable regulatory position is seen in the response in Korea: “the parties must negotiate in good faith. We will not dictate the outcomes of such negotiations, but we will insist that it is undertaken in a genuine spirit of reaching some mutually tenable common position as an outcome.” But what is a “good faith” negotiation if one of the parties maintains that the outcome is unfair? If such negotiations take the form of a set of bilateral negotiations, then presumably a large access ISP with a significant customer population can leverage a more advantageous outcome than a smaller ISP. Does this provide yet more barriers to entry for smaller ISPs with the consequent erosion of a diverse competitive access ISP environment. Do similar considerations apply in the content service space? Are larger content providers able to force outcomes that are more advantageous to them, as compared to smaller content providers? Again, does this provide unwelcome incentives for more aggregation in the content service sector?

In any case, the pressure for some further response appears to be very high in Europe. In a recent media report (<https://totaltele.com/european-commission-to-ask-telcos-for-proof-that-us-tech-giants-should-pay-for-traffic/>) it was noted that there is some enthusiasm for the European Commission to move ahead with such a structure of negotiated termination fees, while other commentators are concerned that there are some serious longer term concerns with the erosion of network neutrality and the potential for network capture as a consequence.

## **Be careful what you wish for**

What happens when the content platforms consolidate further such that the few enterprises remaining in the field are truly massive, far larger than even today?

What happens when these content behemoths have sufficient capital resources to play a dominant role in the access provider by virtue of being their largest and most significant customer through such termination fees? For any enterprise if a small subset of customers dominates the revenue profile for the enterprise, then the enterprise will naturally focus on the needs of these customers at the expense of all others. Ultimately, such a skewed situation could reach the point where the dominant customers effectively take over the enterprise.

This is not idle speculation in this industry. The last decade has seen a comprehensive change in the submarine cable landscape and these days the majority of new cable projects and most of the portfolio of intercontinental cable connectivity is undertaken only by the very largest of today's content providers. More than 80% of the current trans-Atlantic cable capacity is controlled by content providers, not by the traditional carriage industry. So dominant is the content sector's position in this domain that their control of the submarine cable sector is in itself a de facto monopoly. No single carriage entity can amass the

volume of content to justify the construction of new terabit cable project, nor have carrier consortia been up to the task, and few, if any, of the carriage providers have access to investment funds to pay for such a project. The outcome has been that content industry is no longer a customer in this sector. It is also its own provider for this form of carriage.

What then of the last mile access market? Can it sustain itself as an independent sector? Are there sufficient uses and sufficient competitive interest in the access market to withstand the pressure that could be placed upon it from a largely aggregated and highly lucrative content sector? For example, Google has been an interested low-level player in the US access markets for many years. Google's Fi service and Google Fibre appears to be deliberately positioned as low level "foot-in-the-water" exercises that are exploring the space. What might happen if their level of interest in the access sector were to dramatically increase? Consumers would probably be all in favour of such a move, particularly if it involved a dramatic increase in service quality, reliability, and a massive drop in cost to the consumer. What would be the regulatory justification to resist such a move, assuming that it was confidently expected that it would result in a vastly cheaper and more capable digital infrastructure? We have already seen the shift to content-based aggregation for much of or digital service environment, from mail and messaging to document management and commerce, based on offerings that have slashed prices to consumers and increased to robustness and quality of the offered service. The regulatory response to such fundamental changes in this landscape has been a general silence. If the current efforts to force the content sector to engage with the access network operators are indeed successful, and so successful that the content sector takes a substantial position in not just subsidising but directing broadband access infrastructure activities, then what would be left in the access sector? Would they still remain as a separate sector? Or would the broadband access sector become part of an even larger entrenched monopoly of the current suite of content providers?

With particular reference to the European Community, there is already a visible level of concern with the level of offshore control of the regional DNS infrastructure, as evidenced by the DNS4EU program. It appears to me that the concerns over foreign control of critical infrastructure applies as a far greater strategic concern if the subject of the foreign control was the entire national broadband access infrastructure. In that context, inviting these offshore entities to undertake what could become a major investment in this strategically important local asset strikes me as a triumph of short-term opportunism by a small clique of actors at the expense of a broader common strategic interest for the European region, and others of course.

There is also the concern as to whether this could ever be undone if we later regret it. The last time we saw large scale social upheaval was as an outcome of the industrial age at the end of the nineteenth century. Once the industrial behemoths of the day had consumed their competition and become de facto monopolies, they turned their attention from the present to the future and worked diligently to create a dominant position that would echo through the ensuing decades. General Electric, Westinghouse, JP Morgan, Standard Oil are still with us from this period. I have no doubt the current digital behemoths have similar aspirations of an enduring legacy, and once they convert such termination payments into structural subsidisation of access networks, then they can thereby exert greater levels of controlling interest in this critical infrastructure component.

The old adage applies here of being very careful of what you wish for! We really may not have the tools or even the understanding of how to cope with the outcomes were this comprehensive hegemony of content over carriage to eventuate.

---

## Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

---

## Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

[www.potaroo.net](http://www.potaroo.net)