

Robert Kenny

**The attack on settlement-free peering
and the risk of 'access power' peering**

21 June 2013

Contents

- 1. Executive Summary 2
 - Unfounded case for abandoning settlement-free peering 2
 - Strong case against abandoning settlement-free peering 2
 - Conclusion 3
- 2. Introduction..... 5
- 3. The benefits of peering..... 7
- 4. The case made for a move away from settlement-free peering 9
 - Imbalanced traffic and ‘fairness’ 9
 - Payments necessary to support the existing access network and NGA 18
 - Need to create appropriate incentives for CAPs 25
 - Conclusions re the case for a move away from settlement-free peering 26
- 5. The case against a move away from settlement-free peering..... 28
 - Access charges introduce significant friction and gamesmanship into the market 28
 - Paid peering shifts value to one of the least competitive parts of the market 29
 - Paid peering shifts value to one of most opaque parts of the market 30
 - Widespread paid peering creates very bad incentives 31
 - Consumer choice cannot be relied upon to ‘police’ these issues 34
 - If paid peering becomes the norm & consumers can not police, then regulators may have to 37
 - Conclusion re the case against a move away from settlement-free peering 38
- 6. Conclusion 41

1. Executive Summary

Settlement free peering is widespread and plays a critical role for the Internet. By comparison to exchanging traffic via a third-party transit provider, it brings peers significant cost savings, improved performance and greater diversity.

Unfounded case for abandoning settlement-free peering

However, some access providers claim that the near-universal practice of settlement-free peering must be abandoned and content and application providers (CAPs) should pay for peering. Their case is not strong:

Figure 1: Arguments for move to paid peering, and responses

Argument for	Counter arguments
Traffic is imbalanced, with a net flow to access providers	<ul style="list-style-type: none"> • This is true whether traffic arrives at the access provider via transit or peering. A slightly different routing is not a justification for substantial new payments • Net traffic flows tell us little about balance of costs. Access providers transport traffic from a meeting point, but CAPs transport it to that meeting point (often over far greater distances) • Net flows also tell us little about the balance of value. Content and applications are vital for access providers to sell broadband to consumers • Imbalanced costs and value are anyway the norm, not the exception in economic exchange – the critical point is that peering creates value for both parties relative to the alternatives (which it does)
Access providers need funds for traffic growth & NGA	<ul style="list-style-type: none"> • Access providers’ traffic capex is likely falling, as cost efficiencies outstrip traffic growth • Consumers are also spending more for higher speed and higher cap products as their usage grows • Any contribution to NGA costs from paid peering will be trivial
CAPs need cost incentives	<ul style="list-style-type: none"> • CAPs already have substantial costs for distribution, giving them strong incentives to be efficient

Strong case against abandoning settlement-free peering

While the case for shifting away from settlement-free is weak, there are also strong arguments against a general move to paid peering.

Introduction of friction and gamesmanship

Settlement free peering has enabled thousands of interconnects to be put in place on the basis of ‘handshakes’. This brings speed, low transaction costs and flexibility. If such interconnects are to be put on a paid basis, these benefits will be lost.

Value shifted to least competitive & most opaque part of the market

The CAP business is extremely competitive, with much ‘creative destruction’. The consumer broadband access market is reasonably

competitive (in most countries). However, inbound traffic is in effect a monopoly – a CAP wishing to send traffic to a particular consumer has no choice but to deal with that consumer’s access provider.

Paid peering shifts value to this least competitive part of the market, giving access providers a powerful lever to extract monopoly rents from CAPs. Most vulnerable will be national CAPs, (who may be dependent on one or two large ISPs in their own market) and start-ups (who have no balancing leverage from end-user demand).

Paid peering is also opaque – there are no price lists and agreements are often confidential, so monitoring any misbehaviour is difficult.

Paid peering creates bad incentives

One constraint on abusive behaviour is that CAPs may have an alternate route to an access provider – instead of peering they may be able to use transit. However, for precisely this reason, the mere option of imposing paid peering gives access providers a strong incentive to congest their transit links, making them less viable as an alternative. Such congestion harms consumers (by degrading performance for all content relying on transit) and gives access providers the ability to hold CAPs ‘to ransom’. CAPs will be picked off one-by-one, forced to choose between ever more expensive paid peering and congested transit.

Some argue that consumers will make such tactics uneconomic by leaving access providers who congest transit links. However, consumers regularly face performance problems, caused by everything from overloaded sites to interference with their home wifi, and do not necessarily associate them with their ISP (meaning that they may not react to transit congestion by seeking an alternative ISP). Even if they do wish to switch, they face material switching barriers. Thus the consumer response to congestion caused by an access provider may be muted.

Conclusion

The arguments for moving away from settlement-free peering are weak, and there are powerful reasons to be wary of any widespread move to paid peering. It poses a threat to all CAPs, but in particular to start-ups and to national CAPs.

If these players are weakened, or effectively blocked, innovation will suffer, and in turn consumers will be deprived of the creative and constantly improving content and services that have been the driving force of the Internet.

This suggests regulators use their 'soft power' – their ability to persuade and to monitor – to pre-empt any widespread imposition of paid peering by access networks, and in particular any abuse.

2. Introduction

Settlement-free peering is the exchange of traffic between two networks without cash changing hands. The traffic exchanged is generally limited to that between the customers of the two networks. (Interconnect that exchanges traffic beyond this, for instance providing onward delivery to a network's peers and *their* customers, is known as transit).

Historically peering agreements were between entities that looked roughly similar, with each serving a mix of 'traffic sources' (such as websites) and 'traffic sinks', primarily consumers who in general download much more than they upload. As a consequence, approximately equal volumes of traffic flowed in each direction across the peering links.

However, Internet participants have increasingly split between 'access networks' and content providers. Access networks (also known as eyeball networks) are the ISPs that mainly serve consumers and general businesses, that are, as noted above, traffic sinks. Conversely, some Internet participants have focused on offering connectivity to content providers. Further, some larger content providers have developed their own substantial networks.

This bifurcation has led to imbalanced peering links, with substantially more traffic flowing to access networks than from them. According to some of these access networks, this necessitates a move away from the tradition of settlement-free peering.

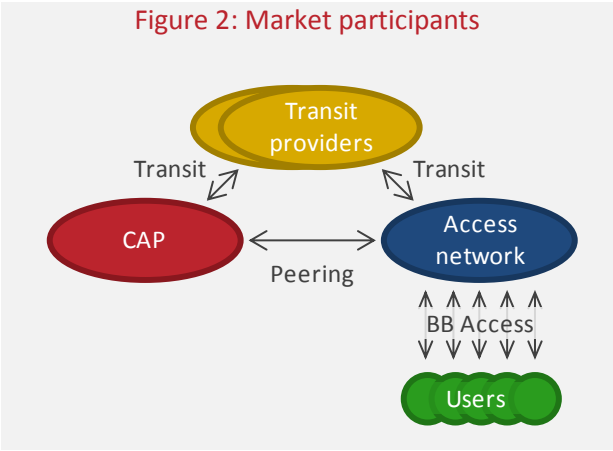
However, such access networks possess a powerful asset. The only way to reach a given Internet user is via his or her ISP. This is known as the terminating monopoly, and is a source of significant market power in setting terms of interconnect. Thus there is a real risk that moving away from settlement-free peering leads to 'access power peering',¹ whereby peering becomes a tool for ISPs to leverage their terminating monopoly inappropriately.

This paper considers peering in the above context, and the trade-offs in moving away from settlement-free-peering to paid peering. To be sure, today some paid peering exists. ISPs may negotiate to get compensation for the costs involved in setting up a peering point (e.g., payment for securing space and power) and sometimes may try to extract additional fees. But, today, this is relatively infrequent and has not so far been treated by access networks as a profit center.

¹ A term coined by William Norton: William Norton, "[Access Power Peering](#)", *Dr Peering Blog*, 6 September 2011

What this paper focuses on, instead, is ISPs' call for a more significant shift away from the norm of settlement-free peering - they are seeking substantial fees and making payment the predominant norm.

In this paper we will (largely) adopt the terminology of the recent BEREC report on IP interconnect.² BEREC referred to 'CAPs' (Content and Application Providers), 'Users' (meaning both consumers and general business users of Internet access) and ISPs. We will split this last into access networks (those directly serving users) and transit providers (those interconnecting other networks), since the dynamics of the two are quite distinct.³



² BEREC, *An assessment of IP interconnection in the context of Net Neutrality*, December 2012

³ Note that some entities may offer both transit and access services, and this overlap will likely be increasingly significant in future

3. The benefits of peering

Today transit represents the ‘default’ way for access providers and CAPs to reach the rest of the Internet. By purchasing from a transit provider, the access provider or CAP gains the ability to reach the entire Internet. A given destination might be reached by going ‘up’ to the transit provider, and then down again to another customer of the transit provider (or a customer of the transit provider’s customer, and so on). Alternatively, if the destination is not within the universe of the transit provider’s customers, the transit provider will pass the traffic on to another transit provider that serves the destination.

While simple and powerful, transit via a third party does have disadvantages:

- It carries a direct cash cost (for both the CAP and the access network), based on volume of traffic
- Compared to a direct connection between two parties, it adds extra ‘hops’ (steps in the transmission of traffic), which add delay
- The longer routing adds more points on the journey where there may be congestion (and hence packet loss)

For these reasons, Internet participants have also chosen to peer – connect directly – rather than exchanging traffic via transit.⁴ This can bring cash savings for CAPs and access networks, and improved performance for both their customers. Traditionally such peering has been ‘settlement-free’, with no money changing hands.

Whether or not to peer with any particular party is a complex decision, depending on whether there is sufficient volume of traffic to justify the fixed cost of a peering link (which will be low if the parties share at least one common location), the technical benefits received and so on. However, there is consensus that peering gives a better experience for end-users. According to Telefónica, for access providers:

“Direct peering is a way to improve the quality perceived by end users and a way to differentiate in the competitive end-user retail market.”⁵

The improvement in perceived quality is of course also valuable to the CAP in question.

⁴ In the early days of the Internet, all interconnect was via peering

⁵ Telefónica, *Public Consultation on specific aspects of transparency, traffic management and switching in an Open Internet*, 15 October 2012

Because of their advantages, peering agreements have become extremely common. The 2011 Packet Clearing House survey (which was certainly not exhaustive) identified 142,210 such agreements worldwide.⁶ This represented 33 per entity, and as such the number of peering agreements is far greater than the number of transit agreements. (The average regional ISP has four transit providers).⁷ That said, transit still plays a valuable role - for “default” connectivity to any network that is not peered, for backup and overflow connectivity, and so on.

The vast majority of peering agreements (99.7% according to Packet Clearing House) are settlement-free, with no cash changing hands. This has been possible precisely because peering is very often beneficial to both parties even without transfer payments.

The convention of settlement-free peering has allowed these arrangements to be based very largely on handshakes. There are widespread common expectations of what parties to settlement-free peering are to provide. This, combined with the absence of cash exchange has obviated the need for formal contracts. The Packet Clearing House survey found that 99.5% of peering agreements were ‘handshake’ based. Settlement-free peering therefore has led to low-cost, efficient transactions between CAPs and access networks.

⁶ Dennis Weller & Bill Woodcock, *Internet Traffic Exchange*, OECD, 29 January 2013

⁷ Amogh Dhamdhere & Constantine Dovrolis, *Twelve Years in the Evolution of the Internet Ecosystem*, October 2011

4. The case made for a move away from settlement-free peering

While settlement-free peering has long been widespread (if not quite ubiquitous), some access networks are now increasingly seeking to emphasise paid peering. Many carriers have publicly made the case for funds to flow from CAPs to access networks,⁸ and likely many have sought such payments from specific counterparties (though generally this will remain confidential).

In a small number of cases, tensions over paid peering have risen to such a point that they have spilled over into the public domain – two familiar examples being France Télécom’s dispute with Cogent and Comcast’s with Level 3.

Access networks’ case for paid peering generally rests on three contentions:

- Traffic flows are increasingly imbalanced, with access networks receiving substantially more traffic than they send – this inherently negates the basis of settlement-free peering
- Payments from CAPs are necessary to support ongoing growth in network capacity, and in particular next generation access (superfast broadband)
- CAPs need to face economic incentives to manage their traffic efficiently

In this section we consider these arguments in turn. We conclude that they are very weak, and certainly not strong enough to justify a radical change to norms that are allowing the Internet to function well today.

Imbalanced traffic and ‘fairness’

Many access networks have cited a shift to imbalanced traffic ratios as a reason to move away from settlement-free peering – in their view imbalanced traffic implies unequal costs, which need to be ‘trued up’ via cash settlement. According to Orange:

“With the increase in asymmetric traffic flows, [settlement-free peering] no longer applies because one party incurs more costs than the other and balance is not guaranteed.”⁹

According to Telekom Austria:

⁸ See for instance AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010

⁹ Orange, *Contribution of Orange France Telecom Group to BEREC Public Consultation On an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

“one reason for having requirements on traffic ratios [for peering] is to be able to agree on a cost sharing of these costs”.¹⁰

ETNO brings value into the picture, quoting Christopher Yoo:

“peering is better understood as a form of barter ... when value is no longer equal on both sides of the transaction, barter no longer makes sense”.¹¹

(In practice many commentators blur the distinction between balanced costs and balanced value, though of course the two are very different).

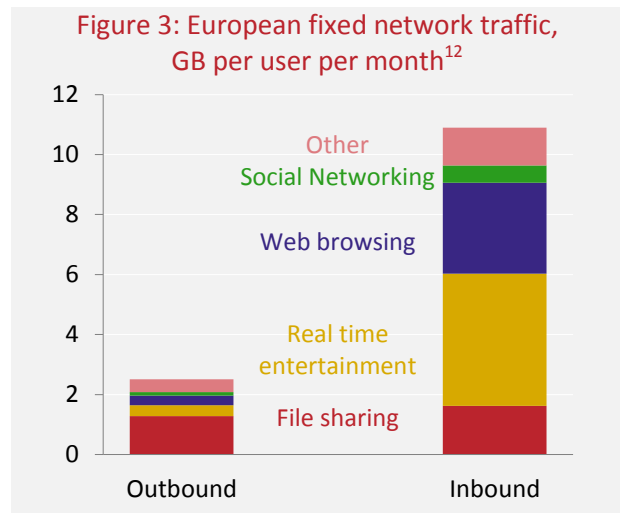
In this section we consider the background to the traffic imbalance; the accuracy of the claim that implies a cost or value imbalance; and its significance if it were indeed true.

Are traffic imbalances increasing?

There is no question that ISP traffic is imbalanced. As Figure 3 shows, for the average European fixed network the ratio of downstream to upstream traffic is now 4.4 to 1 (up from 3.8 to 1 two years earlier).¹³ This does suggest that in:out ratios are increasing, although clearly substantially imbalanced traffic is not a new phenomenon.

Traffic imbalance is to be expected. For both real-time entertainment (Netflix, YouTube, iPlayer) and for browsing, a consumer primarily sends mouse clicks and key strokes up to the cloud, and in return receives far heavier video and images. Further, the products access providers sell encourage imbalance – they generally have far greater downstream than upstream speeds.

This all underlines the significance of how we treat interconnects with a traffic imbalance, since virtually *all* interconnects (be they transit or peering links) between access networks and other categories of network will be in this category.



¹⁰ Telekom Austria, *Telekom Austria Group's Comments on Guidelines for Quality of Service in the scope of Net Neutrality and Differentiation practices and related competition issues in the scope of Net Neutrality and an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

¹¹ Christopher Yoo, *The Dynamic Internet: How Technology, Users and Business Are Transforming the Network*, American Enterprise Institute, October 2012, quoted in ETNO, *ETNO paper on Contribution to WCIT ITRs Proposal to Address New Internet Ecosystem*, September 2012

¹² Sandvine, *Global Internet Phenomena Report 1H 2013*, 14 May 2013

¹³ Sandvine, *Global Internet Phenomena Report Spring 2011*, 12 May 2011


Implications of imbalance for cost

As we have seen, some access networks presume that an increased imbalance of traffic necessarily means an increased imbalance of cost.

The first point to note is that every extra packet sent from a peering point on to an access network has been delivered by the CAP to that peering point. In other words, increased inbound traffic for the access network triggers increased costs for *both* parties, and thus may not alter the cost balance at all.

Even if we set this aside, and consider only the costs of carrying traffic from an exchange point, there remains the issue of distance. The quantum of traffic does drive cost, but so does the distance that traffic travels.¹⁴ Thus a traffic imbalance need not imply a cost imbalance, if the net-sending party is transporting the traffic further than the net-recipient. It is for this reason that carriers such as Level 3 are moving to 'bit-mile' based peering agreements, which seek balance across the combination of distance and traffic, instead of focusing narrowly on traffic alone.¹⁵

This issue is important in the context of peering with access ISPs, since many CAPs are global and have built substantial international networks to bring traffic to local peering points. Microsoft has a network linking peering points in 56 different cities around the world, including 14 in Europe alone:



¹⁴ See for instance Dennis Weller & Bill Woodcock, *Internet Traffic Exchange*, OECD, 29 January 2013

¹⁵ Rajani Baburajan, "IP Transit Provider Level 3 Signs Bit-Mile Balance-Based Peering Agreement with XO", *TMCnet*, 10 January 2013

¹⁶ Adapted from [PeeringDB](#) data, as of 20 May 2013

Akamai (a leading CDN) has a similarly widespread network, with peering locations in 52 cities, in addition to its widespread servers hosted within ISPs around the world:

Figure 5: Akamai peering points¹⁷



Google peers in 59 locations on six continents:

Figure 6: Google peering points¹⁸



This suggests that even if the bits are heavier on the access provider side, the miles may be heavier on the CAP side. (Cost per mile will also be a factor - this may be heavier in the middle mile than in the backbone). The larger CAPs invest substantially to bring content to

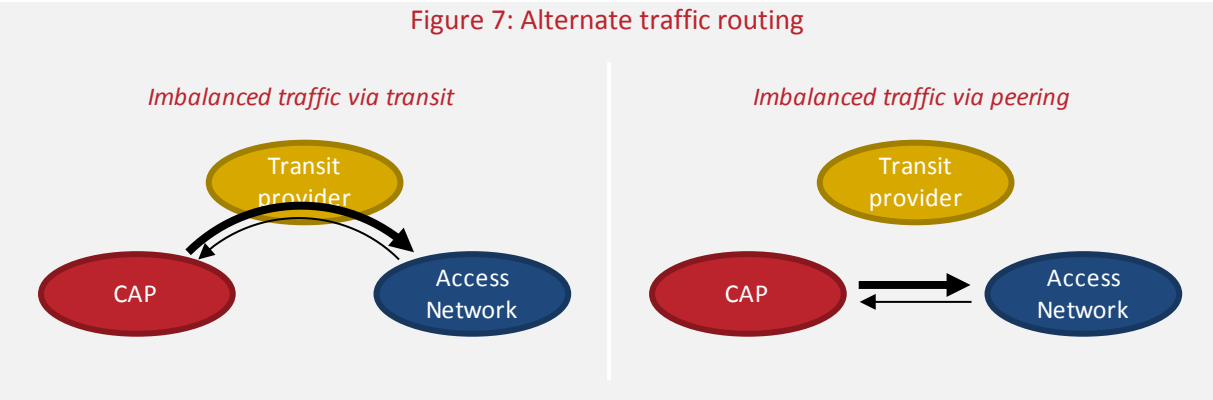
¹⁷ Adapted from [PeeringDB](#) data, as of 17 June 2013

¹⁸ *ibid*

local peering points. In other words, even if traffic is not in balance, this does not imply traffic costs are not in balance.

Even if costs were demonstrated to be imbalanced, there remains the question of relevance. All sorts of economic exchange takes place without the costs of the parties being identical. In the business-to-business context, subsea cable planning, technical standards bodies and airline alliances would all be examples where there is no particular expectation that the costs borne by every party will be identical. Provided that a transaction is beneficial to both parties, it is desirable, regardless of equivalence or non-equivalence of costs.

A final point regarding costs stemming from imbalanced traffic is that they are a result of fundamental traffic flows, *not the peering itself*. If there is a net traffic flow from a CAP to an access provider's users, then this will be unchanged whether that traffic flows via peering or via transit (see Figure 7). Regardless of routing, the access network will face the same imbalance.



If the traffic is identical in either case, it is not clear why the access provider is suddenly entitled to compensation if it arrives via peering, when they receive no such compensation when it arrives via third-party transit (indeed, in that context they will *pay* to receive the traffic).

Implication of imbalance for value

As we have seen, some have instead argued not from cost but from value, presuming that imbalanced traffic means imbalanced value.¹⁹ As with costs, this argument is doubtful both on the basis of accuracy and relevance.

¹⁹ For a more sophisticated approach to determining value splits in peering agreements, see Amogh, Dhamdhare Constantine Dovrolis & Pierre Francois, *A Value-based Framework for Internet Peering Agreements*, October 2010

The idea that net value is in some way simply proportionate to net traffic stands up to little scrutiny. Firstly, the value to CAPs is not strongly related to traffic flow. Consider the following examples:

- Imagine two content providers, one using P2P distribution,²⁰ the other serving centrally. These two providers generate very different traffic flows (with the latter sending much more traffic to the access network via transit or peering). However, all else being equal, there is no reason to believe the revenue of the two would be that different. Thus traffic levels are different, but value to the CAPs is the same
- Conversely, imagine two sites both streaming a similar amount of video to consumers, but one offering video user-generated-content²¹, and the other Hollywood movies. In this case the traffic is the same but the value to the CAPs is very different

Similar issues apply in trying to relate value to the access network from traffic flows. For instance, consider two CAPs sending similar levels of traffic to an ISP. One is a little known video start-up, the other is a leading online encyclopedia much valued by consumers. Clearly the absence of the latter would be much more damaging to the access network's proposition than the absence of the former, and thus though traffic is equal, the leading brand is creating more value for the access network.

The presumption that carrying extra traffic for your counterparty necessarily creates value for them, not you, is a hold-over from the world of voice. If a Telecom Italia customer pays to call a Deutsche Telekom customer then (without an interconnect charge) Telecom Italia is clearly capturing value and Deutsche Telekom is not. The phone network is primarily a 'push' network, in which consumers make calls *to* someone.

By contrast, the internet is primarily a 'pull' network, in which consumers primarily request content *from* someone. Consider a CAP sending requested traffic to Deutsche Telekom broadband customer. The reason that customer is paying Deutsche Telekom for broadband is not for the ability to send requests to the CAP, but to receive the content back. In other words, broadband subscribers are paying primarily to *receive* traffic. (How much would a broadband subscriber pay for a service that allowed them to send but not receive emails,

²⁰ Providers such as Spotify use peer-to-peer to serve their content – that is, customers pull video content not just from a central server, but also from other customers who have already consumed the content in question (and have it stored on their computers)

²¹ Such as consumer uploads on YouTube

request but not see webpages and videos?) By extension, the more traffic received, the greater the value of Deutsche Telekom's product.²²

Indeed, if financial flows are to be related to some notional balance of value, it is far from clear that this results in payments *to* access providers. For instance, if Facebook was in dispute with a small ISP, would Facebook agree to pay the ISP, or might the ISP pay Facebook to prevent a situation where that ISP's customers lost Facebook access? Indeed, this is exactly the situation with much TV content – pay TV operators spend substantial and increasing sums to be able to bring desirable content to their customers.

Thus the balance of value between CAP and access network has little relationship to the balance of traffic. Moreover, there is once again the question of relevance. Even if value was imbalanced, why is this fundamentally a problem? It is absolutely not a requirement of healthy economic exchange that both parties benefit equally. Consider a pharmacy selling aspirin to two customers. One needs the pills to deal with a minor headache. The other needs it urgently to respond to the early signs of a heart attack. Clearly the value to each customer is very different, but the profit for the pharmacy is identical in each case. At least one of these transactions must have an imbalance of value, but does this mean price changes are needed?

Healthy economic exchange *does* depend on the exchange being value-creating for both parties, but settlement-free peering does create value for access providers compared to the default of transit. The strong evidence for this is that access networks always have the option to simply terminate settlement-free peering relationships and revert to exchanging traffic via transit – the fact that they do not suggests that they regard peering as value creating, even without a financial transfer.

Bulk of value not created by peering

There is a final, critical point about the value of peering to CAPs and access providers. The traffic that travels over a peering link creates value for both parties. For the access provider, it enhances the attractiveness of broadband, encouraging adoption and user upgrades to higher speeds.²³ For the CAP, it may bring advertising or other revenue. The vital point is though that both parties can receive substantially the same value if the traffic instead flows via

²² We address the monetisation of this value at page 20

²³ See Page 21 for a more detailed discussion

uncongested transit links (albeit the improved technical performance via peering may bring slightly higher value for both parties).

What this means is that the great majority of value for both parties is not created by peering – rather it is created by *connectivity* (be that via peering or transit).

It would be a radical departure from both the norms of the Internet and telecoms regulation that an access network be in a position to refuse connectivity entirely.²⁴ Indeed, the reason there has been such a strong imperative for universal connectivity is that it creates enormous value for all participants (and certainly it is what consumers expect from an Internet access product is the ability to reach any legal site).

So, if CAPs are ‘entitled’ to connect via transit, then the *incremental* value for them associated with peering is primarily that stemming from the savings in transit costs. However, both parties receive such savings²⁵ - traffic routed via peering represents a reduction in transit costs for the ISP too. Indeed, the access provider may receive a larger saving if it is paying a higher price for transit than (say) a large CAP.

As the traffic in question grows (possibly in an imbalanced manner), the savings grow for both parties equally. In other words, the incremental value of peering is (roughly) equal for both parties, and will remain so even if traffic becomes unbalanced.

Value capture

Of course, this leaves the argument that CAPs (possibly) derive more value from their use of telecoms infrastructure than telcos derive from providing that infrastructure, and somehow that entitles access providers to a transfer payment. Consider ETNO’s request that:

“Adequate commercial arrangements – for example compensation based on specific quality of service or the *value of the traffic* - should be facilitated.” [emphasis added]²⁶

Both the incumbent-funded AT Kearney study²⁷ and a more recent report by Arthur D Little for the Fédération Française des Telecoms²⁸ have emphasised the value capture by CAPs (which they purport to

²⁴ That is, refusing to provide connectivity both via direct and indirect routes

²⁵ Except in the unusual case where the access provider is also a Tier-1 transit provider

²⁶ ETNO, *ETNO response to the Commission Public Consultation on specific aspects of transparency, traffic management and switching in an open Internet*, October 2012

²⁷ AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010

²⁸ Fédération Française des Telecoms / Arthur D Little, *Survey on the “French Telecom Economics” 2012 Edition*, November 2012

demonstrate by showing that certain selected CAPs have had better growth in market capitalisation).

There are numerous problems with this argument.²⁹ Firstly, it is far from clear that any such value capture exists. For example, by focusing on the market capitalization of CAPs that are large today, AD Little and AT Kearney both completely ignore the offsetting loss associated with CAPs that were large in the past (such as Excite, Lycos, Bebo and so on) but which have virtually disappeared today. This is equivalent to assessing the benefits of playing the lottery by only speaking to lottery winners.

Secondly, different value growth between two sectors is not the same as value capture by one of the other. For instance, huge value has been created by search engines, but it is hard to see how this is 'captured' from telcos – it was never a business they were in previously. Indeed search, as a critical enabler of the Internet, has underpinned the growth of broadband revenues for telcos.

Thirdly, even if there was 'value capture' in some way, why should that be a problem that needs remedying? There is no reason whatsoever to expect different links in a value chain to all be making matching returns. As both the AD Little and AT Kearney reports acknowledge, telcos have seen appreciably better performance than the equipment vendors who supply them. If the telcos believe better (purported) returns for CAPs justified a transfer to telcos, should they not also believe that an onward transfer to struggling equipment vendors was also required?

Conclusions re traffic balance

As we have seen, the net direction of traffic flow tells us little:

- It doesn't tell us which party incurs more cost as a result of transporting traffic
- Even if it did, equal costs is not a required or even typical basis for economic exchange
- Traffic flow also tells us very little about who gets greater value out of the exchange
- And again, even if it did, balanced value is also neither required nor typical
- Above all this, the value of the exchange of traffic is created by interconnect, *not* by peering specifically. Unless telcos are planning to refuse to interconnect entirely, the only

²⁹ See also a longer treatment of these issues in Robert Kenny, [*Are traffic charges needed to avert a coming capex catastrophe?*](#), 14 August 2011

appropriate way to assess peering is to look at its incremental benefits relative to transit, which are appreciable for both access providers and CAPs (and may well be greater for the former)

- The overall comparison of value of telecoms and the CAP business is simply irrelevant, as irrelevant as a similar comparison between telcos and their equipment suppliers

It is perhaps for these reasons that those ‘at the coal face’ within ISPs are deeply sceptical about the significance of traffic ratios. According to William Norton (‘ Dr Peering’):

“The Peering Coordinator Community put on a debate on the rationality of peering ratios as a peering discriminator at [industry conference] NANOG 35 in Los Angeles. During that debate, and during the subsequent informal debates afterwards, the consensus was that this metric was neither technically sound nor business rational.”³⁰

Payments necessary to support the existing access network and NGA

Access networks make two arguments from financial need for paid peering and other access charges: firstly that additional funds are vital to support existing traffic growth, and secondly that they are also needed to fund superfast broadband. In at least some instances access networks have been supported by politicians in this regard - France’s Digital Economy Minister Fleur Pellerin recently said “We need to ask serious questions about how web companies can put some money into networks.”³¹

Does traffic growth create a need for additional funds?

Several incumbent telcos argue that traffic growth makes continuation of the status quo commercially impossible – increasing capex to carry additional traffic will make broadband access an unprofitable product and lead to the collapse of the Internet. Telefónica, for instance, has said:

“technological evolution may not compensate [for] the increase of the traffic in all cases. Actually, there are analyses estimating that the costs of upgrading the network capacity

³⁰ DrPeering, *The Folly of Peering Ratios (as a Peering Candidate Discriminator)*. This article is undated, but NANOG 35 was held October 2005

³¹ Marie Mawad, *“France Considers Charging Google for Network Capacity”*, *Bloomberg*, 7 January 2013

are higher than the economies of scale that [are] generated”.³²

According to Telecom Italia:

“it is quite hard to imagine the possibility for telecom operators to bear the burden of the investments absolutely needed to build new networks (or improve the existing ones) in order to cope with the increasing levels of traffic. Lacking investments, the present networks (even though well performing for the present level of data traffic) will soon collapse”.³³

There are several reasons to be cautious about such claims. One is that they are not new – disaster has supposedly been imminent for some time. The analysis cited by Telefonica is a 2010 study by AT Kearney,³⁴ who then said:

“Recent traffic growth figures and mid-term forecasts for future growth ... raise serious challenges regarding the viability of the current Internet model.”

AT Kearney based this view on forecasts that they developed, which indicated that European operator capex for traffic would rise from €15.1bn in 2009 to €28.4bn in 2012, and continue increasing thereafter.

It’s not clear exactly which operators were included in AT Kearney’s analysis, but the European capex (traffic and other) for the EU5 incumbents actually *fell*, from €15.2bn in 2009 to €14.1bn in 2012.³⁵ While this is not an exact match in scope for AT Kearney’s figure, it is very hard to reconcile such a decline with the \$13.3bn *increase* they projected.

There are two reasons why this forecast was likely wrong, and why there is no reason to be anxious about the future sustainability of the broadband business. Firstly, traffic growth has actually declined to relatively moderate rates. For instance, Telecom Italia’s retail broadband traffic in Italy grew just 14% last year, down from 18% the

³² Telefónica, *Telefónica comments on BEREC document BoR (12) 31 Differentiation practices and related competition issues in the scope of Net Neutrality*, 31 July 2012

³³ Telecom Italia, *Telecom Italia response to BEREC Consultation on “An assessment of IP-interconnection in the context of Net Neutrality”*, 31 July 2012

³⁴ AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010, The same study is cited in Orange, *Contribution of Orange France Telecom Group to BEREC Public Consultation On an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

³⁵ BT (total), Telefonica (Europe capex only), Telecom Italia (Italy only), Deutsche Telekom (Europe only), France Telecom (France only). Figures from companies’ annual reports

previous year.³⁶ This compares to AT Kearney’s forecast of sustained growth of 35% per year.

Secondly, growth in traffic is more than offset by declines in unit costs of the relevant equipment. AT Kearney used a highly conservative figure of 15% year-on-year improvement –a more realistic figure is likely twice this.³⁷ However, even if their figure were right, it is larger than the Telecom Italia traffic growth, suggesting that costs will *fall* year-on-year. This seems inconsistent with Telecom Italia’s position that “lacking investments, the present networks ... will soon collapse”.

Users as funders of traffic growth

Even if traffic growth creates a need for additional funding, this does not imply that the money must be extracted from CAPs, since users’ subscriptions are in part directed to traffic costs. BEREC’s view has been that users “incur expenses as they buy network access ... These payments cover both upstream and downstream transmission of data”.³⁸

Access networks argue that extra traffic brings no extra cash from users, since most ISPs do not charge directly for that traffic, and certainly unlimited data caps are common.³⁹ Orange goes further and says:

“Asymmetry cannot be billed in full to end-users, who in general have very little understanding of or control over the traffic generated by their requests”.⁴⁰

This is a puzzling assertion, since the exact same end-users face far more restrictive and complex usage pricing for mobile data services. Orange and other access providers evidently believe such pricing is appropriate in a mobile context, so it is not clear why it would be inapplicable to fixed broadband.

Setting that issue aside, the more general point is that just because the consumer is not explicitly charged for usage on a variable basis (or even aware of their usage) does not mean the consumer is not charged at all. To take a parallel, consider restaurant charges.

³⁶ Telecom Italia, *Annual Report 2012*, March 2013

³⁷ For a more detailed discussion see Robert Kenny, *Are traffic charges needed to avert a coming capex catastrophe?*, 14 August 2011 and BEREC, *An assessment of IP interconnection in the context of Net Neutrality*, December 2012

³⁸ BEREC, *An assessment of IP interconnection in the context of Net Neutrality*, December 2012

³⁹ There are exceptions – they remain a feature of UK and Portuguese pricing, for instance

⁴⁰ Orange, *Contribution of Orange France Telecom Group to BEREC Public Consultation On an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

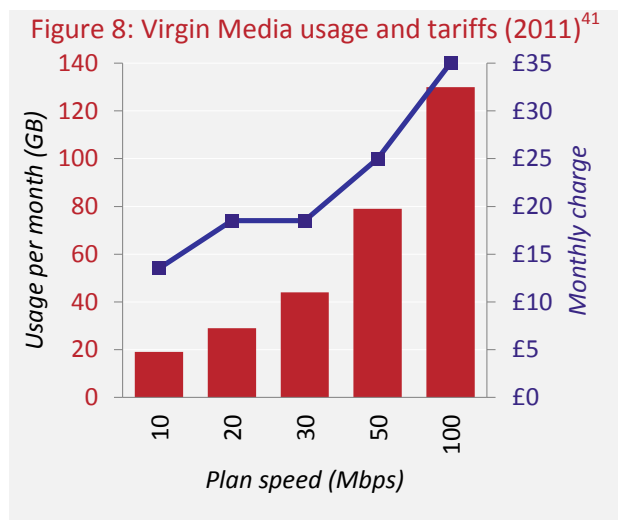
Restaurants do not charge a per-hour tariff for use of the table, or for condiments, or for table linen. However, it would clearly be false to say that the cost of these items was not recovered through the menu prices. Equally, the fact that consumers are completely ignorant of restaurants' property costs does not stop those costs from being recovered (any more than ignorance of traffic levels prevents their cost being recovered from broadband customers by access providers).

Moreover, if the cost of one of these input prices rose over time, this would not require it to be explicitly priced. To extend the parallel, were property prices to rise, that would feed into what restaurants charged their diners, even without an explicit 'table hours' charge.

In fact, access providers do not even need to rely on such indirect mechanisms to extract traffic-related value from consumers. They derive increased revenue from growing traffic, because heavier users tend to be buying higher-speed, more expensive broadband, as Figure 8 shows (using the Virgin Media in the UK as an example).

The direction of causality is not clear. Do heavier users upgrade their speed because of their usage? Or does the capacity of a faster connection encourage people to use the Internet more heavily? For our purposes it doesn't matter. The key point is that additional traffic is very likely associated with additional revenue from users in the form of the premium for a faster connection.

Nor are access providers passive in this – as Figure 9 shows, they encourage users of traffic-heavy entertainment services to purchase more expensive, higher bandwidth products.



⁴¹ Virgin Media, *Q3 2011 Earnings Presentation*, 27 October 2011; Virgin Media, *Q2 2011 Earnings Presentation*, 5 August 2011. Ofcom reports a similar pattern for the UK as a whole – see Ofcom, *Infrastructure Report 2012 Update*, 16 November 2012

Figure 9: Broadband product descriptions for Stofa (Danish ISP)⁴²

Download Speed	Monthly Charge	Description
150 Mbps	419kr	"Lightning-fast up-and downloads of large movies, watch movies in crisp HD."
60 Mbps	349kr	"The whole family can watch TV, movies and play simultaneously"
40 Mbps	269kr	"More computers online at the same time, watch HD and movies online"
20 Mbps	289kr	"Go to Facebook, send emails, play and download music"

Supporting superfast broadband

Beyond a purported need for funds to sustain the current network, some have argued that funds must be extracted from CAPs to support investment in superfast broadband. For instance, Telekom Austria said (in response to BEREC's net neutrality consultation):

"demand for data services is increasing necessitating substantial investments in the future (i.e. in the roll-out of NGA and LTE networks but also in the upgrade of existing fixed and mobile networks) ...It is ... indispensable to find a modus operandi which balances revenues and investments of network operators"⁴³

According to the FTTH Council Europe's response to the same consultation:

"a fair distribution of the Internet value chain is key for industry actors. This is particularly true for FTTH deployments given the investment that is required and the need to monetize those investments."⁴⁴

The fatal flaw with such arguments is that any plausible traffic charges, such as paid peering, will deliver sums that are trivial by comparison to the costs of NGA. (However, note one important caveat - we assume in the following analysis that content providers have a viable transit alternative, a key assumption to which we will return.)

⁴² Stofa, *Speeds and Prices* [accessed 21 May 2013] (Edited Google Translate translation)

⁴³ Telekom Austria, *Telekom Austria Group's Comments on Guidelines for Quality of Service in the scope of Net Neutrality and Differentiation practices and related competition issues in the scope of Net Neutrality and an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

⁴⁴ FTTH Council Europe, *Response to the BEREC Consultation on 'Differentiation practices and related competition issues in the scope of Net Neutrality' / Response to the BEREC Consultation on 'An assessment of IP-interconnection in the context of Net Neutrality'*, 31 July 2012

ARCEP (the French regulator) has undertaken a useful analysis of the breakdown of a typical ISP's costs, shown in Figure 10. As can be seen, transit costs are relatively trivial, at around €0.10 per subscriber. This is significant, because it guides us to how much revenue (or cost saving) an ISP might hope to capture through paid peering.

Let us estimate that today half of an ISP's traffic from CAPs arrives via transit, and half via settlement-free peering. (We ignore traffic from other ISPs, which is relatively unlikely to move to paid peering). Further, let us take an extreme case, where the ISP is successful in moving all this CAP traffic to paid peering.

The first financial impact is that the ISP saves its transit costs - €0.10 per subscriber – since this traffic now comes through a peering link. (Of course, if the ISP is a peer of the transit provider, there will be no such saving). But the ISP now also has new revenue from the paid peering. What might this be worth?

If the ISP's transit links are uncongested (our critical caveat), then some believe that the cost to an CAP of buying transit puts a natural limit on what an access network can charge for peering.⁴⁷ A CAP being offered paid peering at a given price will consider how that price compares to the price of transit via a transit provider. This is an alternative route to the ISP's customers, and so if the cost of paid peering is materially greater than the cost of transit, the CAP is likely to forgo peering – in other words, if the left-hand CAP in Figure 11 is charged more than €0.10 for peering by the access network, it will prefer routing via the transit provider. (As we will see, this logic breaks down if the ISP's transit links are congested). Note that one implication of this is that if transit is acting as a constraint on paid peering, then paid peering charges should be on the same rapid downward path as transit prices.

Figure 10: Breakdown of typical per-subscriber monthly network costs⁴⁵





Figure 11: Hypothetical shift of all access network's traffic to paid peering⁴⁶



⁴⁵ ARCEP, *Report to Parliament and the Government on Net Neutrality*, September 2012. Note that these figures are for an average peak hour usage of 100 Kbps, which ARCEP estimates is today's level

⁴⁶ Note that for simplicity we are assuming the CAP and the Access Network pay the same rates for transit. Figures are per line equivalents, which would in practice be shared across many CAPs

⁴⁷ See for instance David Clark et al, *Interconnection in the Internet: The Policy Challenge*, 9 August 2011

Thus returning to our scenario above where an ISP has moved all its traffic to paid peering, it is unlikely to be able to charge more for that traffic per unit than it itself started out paying for transit (€0.10 per line, for half of its traffic). This means that the ISP’s hypothetical maximum paid peering revenue across all CAP traffic is €0.20 per line. Combined with its transit saving of €0.10, this gives a total benefit of €0.30 per line per month (see Figure 12).

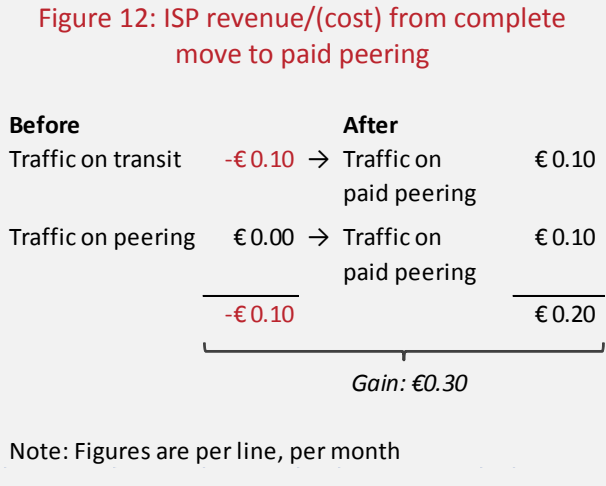
This calculation is extremely generous – in practice no ISP would be able to do away with transit, and we have included the roughly 20% of traffic that is associated with filesharing,⁴⁸ which is unlikely ever to travel via paid peering. Even so, the benefit is trivial in the context of a purported contribution to NGA costs. Analysys Mason estimate the cost per home connected via FTTC at approximately €400 (and FTTH at five times that).⁴⁹ Clearly €0.30 per month would make only a marginal contribution to that cost.

Yves Le Mouel, president of the Fédération française des telecoms, has acknowledged as much, noting that in his view paid peering is:

“only an economic signal providing just a few million Euros, that fill not finance the [superfast broadband] networks of the future, which require investments of €7bn per year”.⁵⁰

Moreover there is no guarantee whatsoever that any extra income from paid peering would be put towards NGA. The €0.30 is not a result of NGA, it is income the operator gets from the transition to paid peering, and consequently it plays no part in an operator’s investment decision regarding NGA.

Thus a move to traffic charges won’t support NGA, and it could in fact be damaging. It is widely acknowledged that one of the material challenges facing NGA is that (as the FTTH Council Europe puts it) there is “no really compelling application yet” that requires it.⁵¹ However, if traffic charges become widespread, it is precisely the kinds of heavy applications for which NGA might be essential that will carry the heaviest financial burden. In other words, traffic charges



⁴⁸ Sandvine, *Global Internet Phenomena Snapshot: 1H 2013: Europe, Fixed Access*, 14 May 2013
⁴⁹ Analysys Mason, *The Italy and Spain NGA cases from a commercial and regulatory point of view*, 19 March 2013
⁵⁰ Guénaél Pépin, “Bercy, le petit théâtre de la neutralité du Net”, *Le Monde*, 15 January 2013
⁵¹ FTTH Council Europe, *Press Conference @FTTH Conference*, 20 February 2013

create a disincentive to develop or offer precisely the kinds of applications that are most needed to create consumer demand for NGA.

Need to create appropriate incentives for CAPs

Orange believes that paid peering is necessary to

“to motivate the [CAP] to better manage its flows and optimise network capacity use”⁵²

Similarly Telecom Italia believes that if content providers are not charged for traffic,

“they will not have any incentive [to become] more efficient and, therefore, occupying the less bandwidth.”⁵³

These arguments miss the vital point that CAPs already have substantial incentives to be efficient with their traffic, in that they pay all the cost to bring it to the edge of the relevant access providers’ networks. Unlike access networks and transit providers, large CAPs generally operate ‘cold potato routing’, transporting traffic on their own networks as far as they can, before handing it over to another network for final delivery. (As we have seen in the example of Microsoft,⁵⁴ some of these networks are massive). This ensures that they carry substantial traffic cost - even if not the full end-to-end cost, since some is borne by users - and have substantial incentives to be efficient with traffic.

Moreover, CAPs without their own global networks often rely on Content Distribution Networks to bring their content to access networks. Globally the CDN market was worth \$1.5 bn in 2012, and is expected to triple in size by 2017.⁵⁵ This is cost incurred by CAPs in delivering their traffic (in addition to their own network costs), and clearly represents a powerful incentive to efficiency.

Looking at Europe in particular, 2012 spending on CDNs was approximately €250m⁵⁶, or roughly €0.15 per fixed broadband line per month.⁵⁷ This one component of content distribution costs

⁵² Orange, *Contribution of Orange France Telecom Group to BEREC Public Consultation On an assessment of IP-interconnection in the context of Net Neutrality*, 31 July 2012

⁵³ Telecom Italia, *Telecom Italia response to BEREC Consultation on “An assessment of IP-interconnection in the context of Net Neutrality”*, 31 July 2012

⁵⁴ See page 9

⁵⁵ Informa, *Content delivery networks: Market dynamics and growth perspectives*, October 2012

⁵⁶ Informa, *ibid*

⁵⁷ Communications Chambers analysis. Note that content delivered via mobile networks remains a relatively small component of total traffic – about 8% according to [StatCounter](#)

already borne by CAPs is only slightly less than the *maximum* paid peering costs of €0.20 per line that access providers might impose on CAPs (again, presuming access providers' transit links are uncongested).

Thus it seems very hard to believe that the incremental incentive to efficiency created by paid peering would be transformative. Moreover, if any access network truly believes it is vital for a CAP to face the incentive of a traffic charge, they have an extremely easy option – they can simply decline to peer with that CAP. As a result, the CAP's traffic would flow through a transit link, with the consequence that the CAP would be paying traffic related charges to the transit provider.

In reality, it is unlikely that an ISP would take such a step, because a consequence would be that they too would face traffic charges to their own transit provider (and they too would face additional incentives to be efficient with their bandwidth). Settlement-free peering is obviously preferable. However, this just highlights the weakness in the access networks' case that there needs to be a substantial move away from settlement-free peering – it is not in fact that settlement-free peering has become unattractive to ISPs, rather that they see an even more attractive option in paid peering.

Finally, if access networks are so concerned that capacity be treated as a scarce resource, it is puzzling that they generally take no steps to encourage their users to be efficient. Clearly pricing to users is at the access providers' discretion (subject to competition), and yet usage caps are relatively rare. As of April 2013, of the pricing plans of the largest ISPs in 19 larger countries in Europe, only 13% included a usage cap.⁵⁸ (This figure is virtually unchanged from 2010, when it was 11%).⁵⁹

Conclusions re the case for a move away from settlement-free peering

As we noted at the beginning of this section, the case for moving away from settlement-free peering relies on three broad arguments: that imbalanced traffic flows are inherently incompatible with settlement-free peering; that access networks need access charges to support growing traffic and the move to superfast broadband; and that CAPs need economic incentives to manage traffic efficiently.

⁵⁸ Communications Chambers analysis of [Google Broadband Pricing Database](#)

⁵⁹ Communications Chambers analysis of pricing data in OECD, [OECD Communications Outlook 2011](#), June 2011

As we have seen, there are severe problems of both fact and relevance with each of these arguments. They are therefore a very weak basis to move away from the widespread, long-standing and successful convention of settlement-free peering.

5. The case against a move away from settlement-free peering

As set out above, the case for a move to paid peering is weak and the benefits are dubious. In this chapter we turn to the costs of such a move.

Access charges introduce significant friction and gamesmanship into the market

The settlement-free framework has made peering simple and quick to arrange. At least 100,000 agreements are in place (and likely far more), with the great majority based on handshakes.

This represents a significant efficiency in the market. Every handshake agreement is time saved in negotiating and contracting those terms. Such agreements also provide fluidity in the market, given the absence of exit clauses.

However, the prospect of adding financial terms will turn each one of these into a battleground, with the two parties debating the relative value received for each such peering agreement. Such debates may be contentious since benefits are both hard to measure (even for the benefiting party), and (because of the heterogeneous nature of the parties) very hard to compare.


Thus friction will be introduced into the market. Man-hours will be lost and the implementation of valuable peering will be delayed. Almost certainly peering agreements that would have been beneficial to both parties will simply not happen, because it is too much trouble to seek terms, or because terms cannot be reached.

Indeed, since peering will become a power play, access providers in particular will have an incentive to ‘walk away’ periodically, in order to demonstrate their seriousness, not just to the particular party in question, but to all access seekers. This will lead to disruption such as that seen in the Cogent/Telia peering dispute, which led to partial outages for Kansas State University, Reuters America, the Swedish Defence Data Agency and many others for two weeks.⁶⁰

⁶⁰ Earl Zmijewski, [“You can’t get there from here”](#), *renesys blog*, 17 March 2008

Further gamesmanship will be introduced since if traffic imbalances are taken as a valid basis to introduce access charges, access networks will have an incentive to create such apparent imbalances. This need not be particularly difficult. For instance, if an access network with a peering link with a particular CAP chooses to route its traffic destined for that CAP via transit rather than via that peering link, then the link will appear unbalanced, since the outbound traffic from the access network has ‘disappeared’ (See Figure 13). Consequently the in:out ratio will rise, giving the access network an excuse to push for paid peering.

Figure 13: Traffic imbalance created by routing outbound traffic via transit



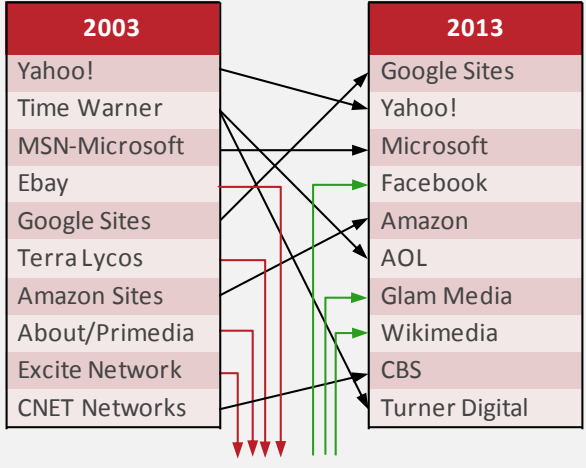
Paid peering shifts value to one of the least competitive parts of the market

In addition to increased friction and transaction costs, an even more significant disadvantage of a move towards paid peering is that it shifts value into one of least competitive areas of the Internet ecosystem, since each access provider has a ‘terminating monopoly’ – that is, anyone wishing to send Internet traffic to a customer of an access provider has no choice but to do so via that provider.

This is quite unlike other aspects of the Internet. In many countries users have a range of access providers to choose from, and competition is fierce in some.⁶² The transit market is more competitive, with a wide choice of Tier-1 and Tier-2 providers.

Consumers also have a vast choice of CAPs, and are willing to exercise it – as the significant volatility in CAP leadership demonstrates. As Figure 14 shows, since 2003 four properties have dropped out of the top 10 (and three of these have essentially disappeared). Conversely, two of the top ten properties in 2013 did not exist in 2003.

Figure 14: Top 10 US websites, 2003 & 2013⁶¹



If the rise of paid peering moves value away from these more competitive domains to the inbound monopoly, then this is


⁶¹ Ranked by monthly visitors. Comscore, *comScore MMX Ranks Top 50 U.S. Web Properties for March 2013*, 1 May 2013 and *comScore MMX Announces Top 50 U.S. Internet Property Rankings for December 2003*, 14 January 2014

⁶² Albeit there are appreciable switching barriers (competition and such barriers are not incompatible)

inherently risky – it gives both the motive and the capability for abuse. (European regulators should be acutely aware of the dangers, having invested significant effort to bring down mobile termination rates.)

While all ISPs have a terminating monopoly, this monopoly is particularly significant for the incumbent operators, since they tend to have substantial market share of the consumers in their home market. Across the EU27, their average share is 43%, and appreciably higher in some markets (Figure 15). This gives them even greater leverage.

Any CAP, large or small, would be reluctant to forgo a particular access provider's customers. However, national CAPs (without presence in multiple markets) are especially vulnerable. It is hard to see how a French CAP (whether start-up or mature) could be viable without access to Orange's customers, for instance.



Such CAPs' situation will be similar to that faced by music websites, which need access to all major music labels' inventory in order to be viable. This gives each label great (and damaging) leverage in their dealings with such sites.

The access providers' leverage means that if a national CAP does not have an alternative, paid peering can be used to extract virtually the entire value created by that CAP, with associated damage to innovation and investment in this sector. This is particularly concerning since a number of access networks are being very explicit both that they feel entitled to much more of the value being created by CAPs, and that they see value based pricing as their route to capturing that value.⁶⁴

Paid peering shifts value to one of most opaque parts of the market

Access networks' pricing to consumers is obviously in the public domain. Tariffs are visible, and regulators and other market participants can easily compare them, both within national markets and between them. This 'sunlight' is a powerful impediment to predatory behaviour and supports the functioning of competition.

⁶³ Communications Chambers analysis of European Union [Digital Agenda Scoreboard key indicators](#). Data is for 2011, except UK which is for 2010

⁶⁴ See page 12

Transit prices are somewhat less transparent in that they are not generally listed on public websites, but they are nonetheless broadly understood. Parties such as Telegeography and DrPeering publish indices of transit prices for multiple cities around the world, a transit buyer is able to get multiple quotes from different suppliers, and so on.

By contrast, paid peering is almost completely opaque. The agreements are bilateral and private, frequently bound by NDAs. Since it may be perceived as weakness to have paid for peering, those who have done so are often reluctant to reveal that they have at all, never mind the price they paid. It is notable from the recent BERC consultation on IP interconnection that there was little consensus between respondents as to whether material paid peering was taking place, and no views were offered at all on pricing levels.

Thus if value shifts to paid peering agreements, it will in practice be incredibly difficult for National Regulatory Authorities (NRAs) to monitor it, and to identify any risks before they become serious.

Moreover, this opacity will make discriminatory pricing incredibly easy – no peering customer will have any idea whether they are paying more or less than other customers of the same network. This will obviously be a powerful temptation for the access network.

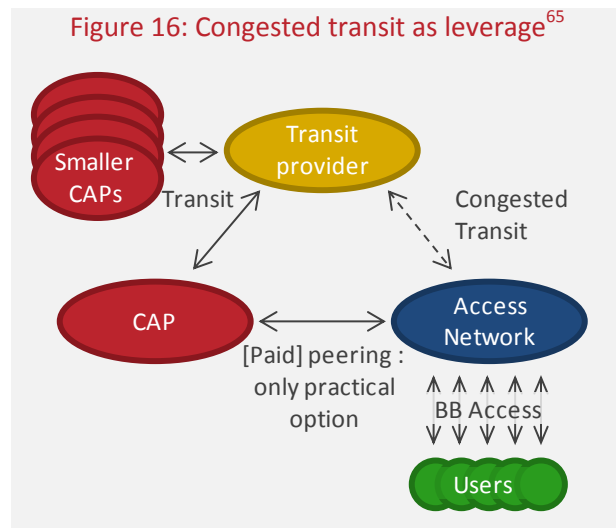
Widespread paid peering creates very bad incentives

Incentive to congest transit links

Thus far in this document, we have assumed that CAPs have an alternative to paid peering – namely, uncongested transit. Such links enable users to reach a global selection of sites and vice versa, creating great value for both and for society at large.

However, if paid peering is ‘normalised’ (rather than being comparatively rare and generally at nominal cost, as today) then access providers will have powerful incentives to do away with uncongested transit. For them, each CAP moved off transit links will be a triple win. The access provider will itself save transit charges; it will see improved performance for users; and it will gain paid peering revenue from the CAP.

This in turn gives the access provider a reason to congest transit links (by not expanding them to meet growing traffic). If transit links are poor, resulting in unreliable service for a CAP's users, then this effectively forces the CAP to move across to paid peering. This is more than a hypothetical – it is strongly believed to be a tool used by Comcast, for example, as they forced open peering links and then converted them to paid peering.⁶⁶ Even more vulnerable to this gambit than the CAPs are the CDNs. As William Norton puts it, if Comcast's transit links are congested:



“From a practical perspective, if you are paid to deliver video to Comcast customers, there is no choice but to purchase Paid Peering from Comcast”.⁶⁷

One particularly unfortunate consequence of this tactic is the ‘collateral damage’. Larger or local CAPs at least have the option of paid peering (inappropriate though it may be). Smaller CAPs, or those distant from the access network, may not be able to justify the fixed cost of setting up a peering connection with the access network, so they are forced through the congested transit link. The poorer performance for these companies (which will include local start-ups) is simply a side effect of the access network using its terminating monopoly leverage to force other players into paying for peering.

Note that if the access network is itself a transit provider, then (local) smaller CAPs have the possibility of buying their transit from the access network, thereby avoiding the congestion. However, this is simply an example of an access network leveraging its terminating access monopoly into the transit market, and indeed gives the access network an additional motive to congest links to other transit providers in the first place. (Such leverage between markets appears to have been a concern of France’s Autorité de la concurrence in the Cogent / France Télécom case).⁶⁸

⁶⁵ The link between the Transit Provider and the Access Network is shown as transit, but the same logic applies if these two parties have a peering relationship

⁶⁶ See Gregory Rose, *The Economics of Internet Interconnection: Insights from the Comcast-Level 3 Peering Dispute*, 28 March 2011 and William Norton, “The Emerging 21st Century Access Power Peering”, *Communications & Strategies*, Q4 2011

⁶⁷ William Norton, *ibid*

⁶⁸ Autorité de la concurrence, *Internet Traffic – Peering Agreements* [Press release], 20 September 2012

Congested transit is also a tactic that inherently builds momentum. With each CAP forced over to paid peering (or even peering), the smaller the set of CAPs remaining on the congested transit link. This potentially allows the access network to degrade the transit link further, since any degradation of this smaller set is less likely to be noted by the network's users. Moreover, from the perspective of the remaining CAPs, ever more of their competitors are on the uncongested peering links, giving them competitive advantage. Thus each CAP that is forced into paid peering makes it harder for the remainder to persist with the transit route. The smallest start-ups (with less traffic and lower financial capacity) are likely to be the last to switch, and thus they (and their potential customers) will suffer most from the poor performance of the transit link.

A situation with congested transit and paid peering is exactly analogous to one with a (poor) 'best-efforts' Internet and premium managed services. The disadvantages of such a two-tier Internet are well known,⁶⁹ and paid peering risks creating this by the back door.

Incentive to create artificial traffic

Paid peering also creates incentives for artificial traffic generation, a problem identified by AT Kearney in their paper⁷⁰ on behalf of four European incumbent carriers⁷¹. In the context of access charges they said:

“a mechanism would have to be put in place to prevent smaller networks fraudulently initiating inbound traffic to generate revenue”.

Clearly paid peering links create exactly the same incentives for an access provider to artificially request content from the peering partner, thereby triggering greater peering revenues. Such tactics have previously been deployed both in the voice world and in the Internet world when there was a gain to be made.⁷²

⁶⁹ See for example: Gary Bachula, *Testimony before the United States Senate Committee on Commerce, Science and Transportation Hearing on Net Neutrality*, 7 February 2006; Lawrence Lessig, *Testimony before the United States Senate Committee on Commerce, Science and Transportation Hearing on Net Neutrality*, 7 February 2006; Robert Kenny, *Are traffic charges needed to avert a coming capex catastrophe?*, 14 August 2011; Bureau Européen des Unions de Consommateurs, *BEREC Guidelines for Quality of Service in the scope of Net Neutrality - BEUC response to the public consultation*, 27 July 2012

⁷⁰ AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010

⁷¹ Deutsche Telekom, France Télécom, Telecom Italia and Telefónica

⁷² For a longer treatment of this issue (and artificial voice traffic in particular) see Robert Kenny, *Are traffic charges needed to avert a coming capex catastrophe?*, 14 August 2011. For artificial Internet traffic, see William Norton, *The Art of Peering : The Peering Playbook*, 2011

Consumer choice cannot be relied upon to ‘police’ these issues

Some access providers have argued that it would not be rational for them to congest transit (or, in a parallel case, degrade ‘best efforts’ Internet) since consumers would switch away in the face of degraded performance.⁷³ For instance, according to Telefónica:

“If an operator reduced best effort to a trickle then customers would not accept this and would switch to a competitor.”⁷⁴

Similarly, ETNO says:

“Due to strong competition in the broadband access markets any network operator that would intentionally degrade the best effort Internet provision ... would lose customers to alternative ISPs and seriously undermine its reputation as a high-quality provider”.⁷⁵

For many regulators, it is the disciplining power of consumer choice that allows them to forebear from regulating directly. For BEREC:

“Competition plays a vital role in guaranteeing net neutrality: the greater the pressure created by competition, the higher the quality of access products and the less incentive an ISP will have to diminish the quality of its own services.”⁷⁶

It is of course incontrovertible that competition between access providers is a powerful, beneficial force. However, it may be expecting too much to presume it can ‘police’ issues such as congested transit routes. There are numerous reasons for this.⁷⁷

⁷³ For an academic articulation of this contention, see for instance Gregory Sidak and David Teece, *“Innovation spillovers and the ‘dirt road’ fallacy: the intellectual bankruptcy of banning optional transactions for enhanced delivery over the internet”*, *Journal of Competition Law and Economics*, 5 July 2010. Note that Sidak and Teece primarily address the US market. Whatever the merits of their case in that context, it is greatly weakened in the European market. One of their key assumptions is that ISPs face low variable costs, so the loss of a subscriber brings substantial loss of profit. In Europe, where many ISPs are dependent on wholesale products (such as unbundled local loops) this is clearly less true

⁷⁴ Telefónica, *Public Consultation on specific aspects of transparency, traffic management and switching in an Open Internet*, 15 October 2012

⁷⁵ ETNO, *ETNO Response to the BEREC Consultations on: Guidelines for Quality of Service in the scope of Net Neutrality – BOR (12) 32 and Differentiation practices and related competition issues in the scope of Net Neutrality – BOR (12) 31*, July 2012

⁷⁶ BEREC, *Summary of BEREC positions on net neutrality*, 3 December 2012

⁷⁷ These issues are discussed in detail in the companion report, Communications Chambers, *Consumer Lock-in for Fixed Broadband*, July 2013

Switching barriers

Consumers face material switching barriers in the face of network degradation. These include (but are certainly no limited to):

- The consumer may simply not notice the degradation, even if it is material, since faults and variations in quality are common (for reasons within their ISP's control and beyond it). Thus any incremental degradation caused by a congested transit route may be 'lost in the noise'
- Even if a consumer notices the degradation, he may not attribute it to his ISP (as opposed to, say, a problem with his computer, home network, the site itself and so on)
- Even if the consumer blames the ISP in response, he may be locked in by contract
- Even if the consumer is able to switch, he may be put off by the time and effort of switching, or he may not feel certain that he is able to identify an alternative ISP that would have better quality
- The consumer may be taking a bundle, and may (for instance) be reluctant to relinquish favourite TV content simply to get a different broadband supplier

It is for these reasons that broadband markets in general are becoming increasingly stable, with less market share movement each year.⁷⁸

Asymmetry between CAPs and users

Even if consumers were not increasingly locked in, the key question would not so much be their absolute level of inertia in the face of degradation, but rather how it compares to that of CAPs. If an access provider is congesting transit to exert pressure on CAPs, but causing some degradation for its own consumers, the critical issue is 'who blinks first'. There are many reasons to think it is likely to be the CAP.

For instance, levels of congestion that are barely detectable by consumers may be very damaging (and very obvious) for CAPs. The latter have the technical capability and the expertise to monitor performance constantly, and are aware of the cash cost of degraded connectivity

A further issue for CAPs is that one access provider is likely far more important to them than one CAP is to consumers (though some of the largest CAPs may be exceptions). For a national CAP, a given access provider could easily represent more than 50% of their target

⁷⁸ Communications Chambers, *ibid*

market. Degradation on the links to that provider could be devastating. Conversely, for the users of that access network, the CAP will be one amongst very many they deal with on the Internet (indeed, many users may not deal with that particular CAP at all).

This lack of switching barriers and the asymmetry between the CAP and access provider players of the ‘game of chicken’ suggest that access providers will be in a powerful position to force CAPs onto paid peering.

Access networks’ ability to pick off CAPs

This is particularly true given that access providers will increasingly be in a position to be quite targeted in their efforts to force CAPs onto paid peering. Consider a major CAP who is currently on settlement-free peering, but the access provider wants to push them to paid peering. The access provider runs their transit link to provide acceptable performance, but with little slack. It then presents the CAP with a requirement to transition to paid peering.

This presents the ‘target’ CAP with the choice of accepting the move to paid peering, or of moving its traffic over to transit (much as did Level 3 in its dispute with Comcast). Particularly for a CAP sending significant traffic, this move itself could meaningfully congest the transit link, causing problems both for itself and any other CAPs using transit.

However all the other peered CAPs are unaffected. Thus even if the target CAP is brave enough to trigger congestion, users will not see any degradation for many of the CAPs they use. One consequence is that those users may be less likely to attribute to their access network any performance problems they happen to notice with the CAPs on transit. By extension those users will be less likely to switch away from the access network.

Once this target CAP is ‘brought to heel’, the access network can then move on to the next one, repeating the process.

As we have noted, the bilateral nature of paid peering makes it very easy for an access network to be highly discriminatory in its pricing, extracting maximum value from those with higher capacity to pay, or from those whose services represent a threat to the access network’s own operations. For instance, an OTT video provider could easily be charged a punitive rate by an access network with its own TV offering.

In this aspect, paid peering is considerably worse than an access network providing a degraded ‘best efforts’ offer and a premium

tariff for guaranteed delivery. At least in the latter case, the premium tariff will be declared and available to all on a non-discriminatory basis.

If paid peering becomes the norm & consumers can not police, then regulators may have to

If consumers are not in a position to ‘police’ paid peering, then regulators may in time be forced to. We do not advocate such intervention, which would carry significant cost and risk – far better that the situation never arose where it became necessary. However, a time could come where it was the ‘least bad’ outcome. To prevent rent seeking by holders of terminating monopolies, regulators might be compelled to intervene in one or more of the following areas:

- Preventing undue congestion in transit links (or, in the case of a vertically integrated Tier 1 operator, preventing congestion in that operator’s peering links with other Tier 1s)
- Ensuring that access charges are not excessive
- Ensuring that access charges are non-discriminatory

As BEREC have noted in the context of traffic classes:

“Traffic classes using prioritisation introduce an incentive to decrease the quality of the “best effort” class vis-à-vis premium classes to create a willingness to pay for premium quality. This creates the need for more regulatory control including the potential need for a minimum quality of service, introducing additional monitoring requirements.”⁷⁹

Exactly the same logic applies to paid peering and (potentially congested) transit. Unfortunately, such monitoring and control will, for the reasons set out above, be particularly difficult given the opacity of paid peering.

A practical example of paid peering necessitating greater regulatory intervention is the recent Cogent / France Télécom case. In considering this dispute, the Autorité de la concurrence approved France Télécom’s move to paid peering with Cogent. However, they noted that there was no formal transfer pricing between FT’s retail and transit businesses, creating the risk of margin squeeze (once settlement-free peering fell away). The Autorité has required FT to formalise transfer pricing and procedures and to set up a system to monitor these.⁸⁰ That said, unless its findings are made public, it will

⁷⁹ BEREC, *An assessment of IP interconnection in the context of Net Neutrality*, December 2012

⁸⁰ Autorité de la concurrence, *Internet Traffic – Peering Agreements* [Press release], 20 September 2012

not enable other market participants to tell if they are being fairly treated by FT.

To reiterate, we are not calling for regulation of peering or transit as a desirable outcome. Such interventions could do more harm than good, since they would remove some of the flexibility and responsiveness that has been such a vital aspect of the Internet's success to date.

This is another reason why finding some way to maintain the standing, predominant norms of settlement-free peering is crucial. Otherwise, there may come a point of no return, where today's norms are gone and a regulatory solution, no matter how inefficient, is needed.

Thus we suggest regulators use their 'soft power' – their ability to persuade and to monitor⁸¹ – to pre-empt any abuse of paid peering by access networks. One of the reasons for suggesting such an approach is that it will reduce the likelihood of NRAs ever having to use their 'hard power', in the form of mandated interconnect, regulated prices and so on.

Conclusion re the case against a move away from settlement-free peering

It is not the contention of this report that paid peering is inherently wrong. In some circumstances it may be justified – for instance, it may suit a small regional network and a large global wholesale provider to connect via paid peering rather than a transit relationship. However, like explosives, paid peering is a dangerous tool that can do great damage in the wrong hands. It:

- Creates an incentive to congest transit links, a hitherto vital component of the Internet
- Gives access networks the potential to leverage their terminating monopoly to extract value from CAPs
- Enables discriminatory pricing
- Shifts value to a domain where it will be very difficult for NRAs to monitor and intervene if necessary

These issues would be less concerning if they were hypothetical. However access networks have already demonstrated their

⁸¹ In the context of peering agreements, monitoring needs a light touch – for instance, a requirement to report the terms of all agreements could require formalisation of those agreements, damaging the current 'hand shake' approach. Limiting reporting to those agreements involving financial payments above a certain threshold would be a way to mitigate this problem

willingness to congest transit links to secure paid peering, and have declared their view that they are entitled to a greater share of the value created by CAPs.

It is possible to see access providers' pursuit of paid peering with CAPs as the current step in a longer term strategy of gradual value capture by access providers. In the early days of the commercial Internet, access providers were dependent on Tier-1 providers (or regional transit providers) for interconnect. This resulted in some highly inefficient 'tromboning' of traffic to the US, and in reaction access providers moved to local peering. This deprived the Tier-1s of revenue, but was clearly economically efficient.

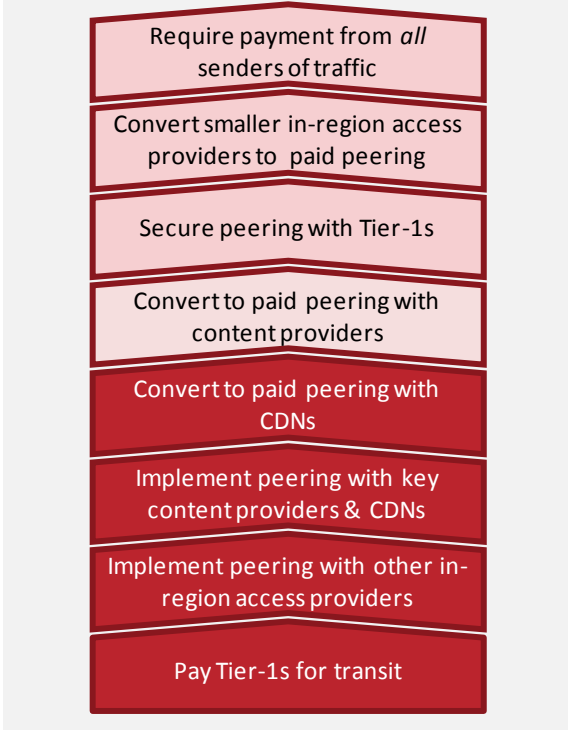
Subsequently both larger CAPs and CDNs started to peer with access providers as well, enabling their mutual traffic to avoid transit expense. Again, this brought clear economic benefits.

The next step was for access providers to insist on payment from CDNs. As noted, CDNs were particularly vulnerable to such demands, since their ability to deliver traffic was fundamental to their own customer relationships, and those customers could readily switch between CDNs.

Now, access providers are pushing to convert CAPs to paid peering. As we have seen, there are many reasons to be very sceptical of such a move on its own merits. However, it is unlikely to be the last step in the game. In just the way in which access providers are seeking to use the leverage of the terminating monopoly to extract value from CAPs, they will move on to doing the same with Tier-1s, initially by moving to peering, not transit, and then by converting those links to paid peering⁸³

Two examples in the public domain where this has happened already are Comcast/Level 3 and France Télécom/Cogent. In each case the access provider had existing settlement-free peering links with the Tier-1 in question. However, on the premise of imbalanced traffic, the access providers required the Tier-1s to pay for interconnect.

Figure 17: Value capture by access providers⁸²



⁸² Illustrative – note that the steps shown need not be taken in strict sequential order

⁸³ See for example Timothy Lee, [“Keeping the Internet Competitive”](#), *National Affairs*, 20 March 2012

In effect, larger access providers are becoming integrated Tier-1 providers, paying transit to no-one and extracting termination fees from an ever wider group of players, including the current, independent Tier-1s and smaller local access providers. Indeed, in this scenario the independent Tier-1s will have the same vulnerability as the CDNs discussed above – they will have promised their customers universal delivery, but in order to fulfill that promise they will have no choice but to pay the fees demanded by the new, integrated Tier-1s.

The end-game is that everyone will pay access providers for the delivery of traffic, with those providers coordinating their prices to prevent any CAP seeking negotiating leverage. This is not a paranoid fantasy – it is what leading access providers have publicly set out as their desired outcome, stating (via AT Kearney):

“a more coordinated approach is probably required where all Retail Connectivity Providers in a market set similar charges”⁸⁴

The question for NRAs and other market participants is whether this is a desirable endgame, and if not, when to draw the line.

⁸⁴ AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010

6. Conclusion

Settlement-free peering has been a fundamental underpinning of the Internet since its earliest days, and has thereby contributed to the Internet's amazing value creation (for consumers, access providers, CAPs and society in general).

This suggests caution in seeking to move away from settlement-free peering, with a high 'burden of proof' for a view that wholesale change is necessary. In fact, as we have seen, the evidence that there is a problem with the status quo looks to be weak. Access providers appear to be adequately compensated (by end users) for the carriage of traffic, and certainly the frequently predicted unsustainable capex explosion has failed to materialise year after year.

Moreover, the 'tactical' basis for a purported need to switch to paid peering – traffic imbalance – is long standing, and has little relevance to the cost or value distribution between parties (even if one were to grant such distribution as a relevant factor to price setting).

If the case for moving away from the current well established, well functioning system is weak, the dangers in the new proposed structure are substantial. If paid peering becomes the norm, the existence of a paid path will give ISPs a strong incentive to degrade the transit path. Such degradation will cause harm, but it is unlikely in practice to lead to significant levels of consumer switching between access providers (which might otherwise act as a disincentive to such degradation).

Moreover, a two-tier Internet based on paid peering and congested transit will be particularly open to the exercise of the power accruing to access networks through their terminating monopoly. Since paid peering is so opaque, regulators will find it deeply challenging to detect or respond to any abuses.

BEREC has taken the view that

“If ... practices [such as] paid peering ...became widespread where Internet Access Provider connecting end-users were able to set abusive charges for interconnection out of a monopoly position, this outcome might not be considered desirable.”

Against the evidence set out in this report, BEREC's view is a masterful under-statement.⁸⁵

⁸⁵ BEREC, *An assessment of IP interconnection in the context of Net Neutrality*, December 2012