Draft report for public consultation (BoR 12-33)
An Assessment of IP-Interconnection in the context of Net Neutrality
Comments from Cogent Communications

Comments Summary

BEREC has rightly identified the separation of network and application layers as the defining characteristic of the internet. It is certainly the single most important technical factor allowing the creation of applications, services and contents on the internet.

However, the distribution of those services is just as essential if their benefits are to be reaped by the whole economy. This requires a communication system that provides "global connectivity", is efficient, seamless and fosters competition between the various methods of interconnections. The internet, built around "best effort" transit and peering agreements is already delivering those benefits.

The Internet is a massively disruptive machine. Incumbent telecom operators are doing their upmost to hinder its progress. BEREC's report on Net Neutrality violations is a stunning indictment of Incumbents' behavior and justifies the toughest of regulatory responses.

Within this context, interconnections placed at the gates of their monopoly network, provide incumbents with the first opportunity to distort competition between "*off-net"* services that are available on the net and their internal, competing "*on-net*" offerings. Their walled gardens are already in place, as traffic requested by their users is already hindered by organized congestions or by financial discrimination. Relying on end users' ability to switch providers is not enough to keep incumbents honest. Regulatory guidance is needed.

"*Off net*" content distribution can be achieved through different competing methods (transit, direct peering, CDN or participation in IXPs). All of them, except for transit, are one-to-one relationships where incumbents' retail market power can be applied. Transit is not only the keystone of the Internet; it is also the only weapon that can insure a proper competitive landscape between all distribution methods.

BEREC correctly identifies that transit needs always to be an option and that peering and transit need to be substitutable. Unfortunately, regulators have never taken any steps to ensure this was effective. They now need to make sure that it is, and need to clarify incumbents' obligations under Art.5 AD. We identify three essential conditions for transit to always be an effective default option. Interconnections must be congestion free, local and end users' prefixes propagation must not be artificially withheld.

Given the importance placed by Regulators on the ability for subscribers to switch internet suppliers, BEREC should also consider one alternative option, i.e., the separation between the physical access network and access to the global internet. This would follow naturally from the segregation of network and application layers. Network owners (i.e., Incumbents) would charge for network connection at the

hardware level, but should be prevented from selling global connectivity bundled with it. Users could then pick whichever connectivity offer best suits their needs. These can be obtained from alternative providers or even from Tier-1 operators directly.

Responses to specific BEREC questions

Question 1: Are any other important players and/or relationships missing?

There are three major relationships missing or needing emphasis. They revolve around traffic causality, structural ratio imbalances and competition.

Causality:

It is a simple fact that, in the Internet world, as opposed to the voice world, traffic is caused in both directions by the same end user. CAUs generate the download requests and those are causing the responses received from CAP's.

Some eyeballs ISPs argue that CAPs send more traffic than necessary when answering the download request. Whether this is right or wrong does obviously not change the causality of the traffic. It only argues about its quantity. It is not for Cogent to comment on CAPs business models, but we believe that they compete on speed, quality of experience and consumer satisfaction. The size of their responses is part of that equation, and if consumers suffer from information overload, they can easily switch to a different CAPs.

It is therefore inaccurate to look at net neutrality violations on the interconnection level as "*a situation where end-users cannot reach all destinations*". The definition must be broader and include the return path as well. It should read "reach and receive traffic to/from all destinations"

Structural ratio unbalance:

It is also a simple fact that each end user determines, by his own behavior, his own specific ratio of incoming and outgoing traffic.

A minority of users may be responsible for the vast majority of the traffic. Eyeball ISPs can influence their behaviors by offering various transparent payment plans. Consumers 'responses to those plans (via switching) will testify to the competitive state of the access market.

At the interconnection level, the incoming / outgoing traffic ratio between Eyeball ISPs and the rest of the internet is naturally / structurally unbalanced. It represents the sum total of all CAU's individual surfing behaviors after the impact of Eyeball ISP's potential transparent payment plans[1].

Some Incumbent ISP's argue that the ratio imbalance between themselves and the rest of the world justifies the imposition of a payment (access tax) from transit providers or CAPs to terminate the traffic requested by their CAU's. Ratios have some use between operators of the same kind and we discuss them further down in this contribution. However, ratios between operators of different kinds have no relevance. Not only is the imbalance is structural, but access networks are unavoidable and control both the numerator (incoming traffic) and the denominator (outgoing traffic) of the equation.

Any limitations or constraints at the interconnection level [2] to this sum total of aggregate traffic amounts to a non-transparent restriction to CAUs freedom to access all applications, services or contents of their choices, and constitutes a violation of Net Neutrality.

BEREC should categorically affirm that subjecting CAU's aggregate internet usage to limited or restricted transit or peering capacity constitutes a violation of the net neutrality principle. [3] In other words, the **first essential condition** that IP Interconnections need to fulfill is to be **congestion free**. This would mandate eyeball ISPs to purchase enough transit to satisfy CAU's download requests, but would also mandate those ISPs that rely on peering-only to dimension those ports without any constraints.

CAUs need to be able to make intelligent and well informed choices. If sunlight is the greatest disinfectant, then transparency in the interconnection area must also be mandated. CAUs have a right to know if their internet freedoms are being impeded by their eyeball ISP at any level.

[1] CDN servers placed inside the Access Provider network will reduce this ratio, but CDNs are only one alternative method of delivering the requested traffic and needs to be in analyzed in a fair competitive environment with other alternatives (Transit, peering and direct interconnections) – see below -

[2] This can happen if Eyeball ISPs fail to purchase enough transit or provide enough peering capacity to accommodate the incoming traffic requested by their CAUs.

3 This affirmation would be in line with the recent judgment from the Tribunal de Grande Instance in Paris, which condemned FREE for not providing enough capacity to satisfy their users' download requests in unbundled areas. This judgment is currently under appeal. (TGI Paris, 31e chambre, aff. 0700396023, 7 février 2012)

Competitive methods of content distribution

Eyeball ISPs satisfy the download requests from their CAUs through a mixture of transit purchases, private settlement free peering, direct paid peering, CDN and participation in public peering exchanges. Each of those methods has its own technical and financial specificities. At the other side of the distribution chain, CAPs must be able to choose independently whichever of those methods best suits their needs. A healthy state of competition must be maintained between all those distribution methods. BEREC and The European Commission have consistently underlined that substitutability between those various methods is critical.

However that healthy state cannot be taken for granted. Every method, except one, mentioned above is a "one-to-one" relationship that only provides limited connectivity to a limited number of networks and CAPs. Transit stands alone among this group as it is the only method that allows all traffic coming from all origins to be received: it is the only method that provides "global connectivity" and can therefore compete with all other methods. BEREC correctly identifies that transit is the "default" option.

But for this default option to be effective, it needs to enjoy the same characteristic as the others.

Most eyeball providers do not have an associated transit business and naturally buy transit locally. However, those that have and rely on peering-only relationships to deliver internet access to CAU's typically insist that those interconnections be located away from their local market. This introduces QoS differences for local traffic (increased latency) and is widely used as a method to force local content providers to peer directly (at a cost) with the eyeball provider.

In other words, for transit to be a true default option, the **second essential condition** that interconnections must fulfill is to be **local,** i.e., provided in the same home market as the direct peering alternatives are –

Question 2: Do you agree with the classifications of CAPs as outlined above?

Question 3: Do you agree with the classifications of CAUs as outlined above?

In general, this is a fairly good description of the ecosystem – CAP and CAU are properly described and constitute the only two players that are net creator / consumer of IP traffic. By opposition, all other categories are intermediary networks or ISP's through which traffic passes.

Question 4: Do you agree with the classifications of ISPs as outlined above?

The use of "ISP" is too general to describe those intermediaries as their roles and market positions differ widely. The description of an "Eyeball ISP" is accurate, and so is the description of the "CAP focused ISP", which are commonly called "Hosters". Yet, there is a critical difference between them, which BEREC could have highlighted further.

Eyeball ISPs are best depicted as a series of side-by-side monopolies since all their clients are naturally single homed. This makes their network unavoidable. Of course, CAUs have the ability to switch providers (even though there are numerous obstacles in doing so and they tend to be quite sticky), but they would still remain single-homed, after the switch, with the competing Eyeball ISP.

BEREC correctly identifies that those ISPs also sell "*services that compete with over the top providers*". This lacks emphasis. Those services are "*on-net*" for the eyeball ISP and the desire to sell them to their CAUs and to discriminate against the competing "*off-net*" services available on the internet is the root cause of all net neutrality violations observed in the interconnection market.

CAP focused ISPs or Hosters operate in an intrinsically competitive market. CAPs are naturally multi homed and purchase connectivity through their hosters (typically part of a bundle comprising also server rental, collocation and power) but also through direct transit purchases, CDNs or direct connectivity with Eyeball ISPs. CAPs are constantly evaluating the relative benefits of the various distribution methods and are modifying their mix of suppliers permanently.

BEREC actually put Cogent into this "CAP focused ISP" category, alongside Atrato and 1&1. This is either a simple oversight or a complete misunderstanding of Cogent's business. Cogent is a pure transit provider and belongs to the third category below.

The third category of ISP needs a much more thorough description, which Cogent would like to provide below

The Backbone or Transit ISP. Without these ISPs, the internet would simply not exist. None of the two previous ISP categories can, by themselves, provide internet access (i.e., the global routing table) to their end clients. Transit Providers are the only source of "global connectivity" which is why they constitute the "keystone" of the internet. Their characteristics and interactions must be highlighted.

1. **Transit.** Transit operators are economic animals. They strive to maximize revenue and minimize costs. This explains how client traffic is handled and how transit providers interact between themselves.

    a) Outgoing traffic - Any network is always in complete control of its outgoing traffic. Every participant (transit ISP, CAP, Hoster or Eyeball ISP) can decide unilaterally to send all or parts of their outgoing traffic to any network that they connect to. For traffic received from a client and destined to a non-client, the need to minimize costs will dictate the nearest exit routing protocol (see hot potato below).

    b) Incoming traffic and route propagation mechanism. How incoming traffic is handled follows the need to maximize revenue. Prefixes (blocks of IP addresses) are announced by the client to its transit provider(s) and those are propagated / advertised widely by the transit provider(s) to its other clients and to its peers (and then onwards to the clients of those

peers) in order to attract as much traffic as possible to those prefixes. This route propagation mechanism constitutes the very heart of the internet and is solely responsible for establishing global connectivity.

BEREC correctly mentions that "*An access provider however is not entitled to any payment when taking over down-stream traffic at his agreed PoI and physically terminating a data flow as this is paid for by the end-user*." This establishes the fundamental principle of Bill and Keep (more on this later) charging method for the internet.

Implicit in any transit contract therefore is a difference in network load impact for traffic that is received from a client and destined to a peer (traffic is routed out as soon as possible to minimize network load) or destined to a client and received from a peer (traffic is accepted anywhere on the network to maximize revenue from the client). We will return to this issue further below in the settlement free peering section.

c) Route Multiplicity: Clients (CAP, hosters and Eyeball ISPs) are naturally multi homed and the same prefixes are therefore propagated through a multitude of alternative paths. The end client can manipulate various functions in the BGP session (AS path prepending mainly) to try and steer incoming traffic through one specific transit network, but the ultimate control (the decision to send traffic through a specific transit provider) rests with the sender of the traffic. The end client can also manipulate another function in the BGP session (called "communities") to prevent route propagation to specific networks. This use of communities can lead to abuse, as will be seen below.

Access providers choose independently where to connect to the internet. While this is usually done within their geographic coverage area, the location(s) of those interconnections is not necessarily neutral in terms of QoS. In particular, one would wonder why an offer to interconnect deeper into the network (at a regional rather than national level) would be refused by large access providers, since this would, at no incremental cost, increase QoS for the user and reduce network load on their internal backbone. That question becomes even more relevant when access providers exclusively purchase transit from a single supplier and when this transit provider only offers peering capacity outside of the access provider's national reach[4].

Technically speaking, those route announcements are copied in the routing tables from one router to the others in real time and provide redundancy to the entire system. If any of those routes were to become unavailable (because of fiber cuts, equipment or power failures or more simply because of a change in transit provider by any party), traffic would still flow automatically via an alternative path. Importantly though, routers do not deal properly with congestions on the preferred path. In case of congestions, IP packets continue

[4] This is specifically the case in Poland, where TPSA buys transit from France Telecom and France Telecom does not provide its peers with interconnections in Poland.

to be sent along the same path and are eventually lost as they drop from the routers' buffer memory, resulting in QoS degradation. IP engineers can manually force traffic through a different path (if there is one) but it always results in a lower quality experience (longer paths and induced latencies) than necessary.

Importantly, the combination of route multiplicity and route propagation mechanism allows for all clients to make transit purchase decisions independently from each other, without having to worry about any loss of global connectivity.

2. **Peering** - By opposition to transit, peering does not allow for route propagation to other peers and cannot therefore offer global connectivity. Connectivity is reduced to the clients of the peered networks. It is a one-to-one relationship that only provides route singularity[5]. Peering is therefore, as properly identified by BEREC a "non-transitive" relationship.

   Technically, peering is the shortest and best quality route between the respective clients. If they have their clients' interests at heart, both networks will want to maximize this traffic and will therefore dimension the peering interconnections accordingly. Within that context, it is acceptable for a client to use "communities" to prevent route propagation from a transit supplier to a network that peers both with the transit supplier and with the end client. However, the use of those communities, in conjunction with congestions on the peering interconnections, leads to the suppression of alternative paths and the ability to exert uncompetitive commercial pressures on the direct "one-to-one" peering relationship.

   This allows us to identify the **third essential condition** that Interconnections must fulfill for transit to be an effective default option: **end users' prefixes propagation must not be withheld artificially**. While this is usually not an issue if interconnections already fulfill the first two essential conditions (i.e., congestion free and local) it becomes a critical one when they are not. Regulators should look at this essential condition as being the transit provider of last resort, or as transit's "safety wheel".

   It is worth mentioning that peering is the most overused word in the internet world[6]. There really are two kinds of peering: settlement free and paid or "direct". Settlement free peering typically exists between networks of the same kind (transit/transit or access/access), while "direct" typically refers to networks of different kinds (CAP/Eyeball) and is usually associated with a payment received for terminated traffic. From and eyeball ISP perspective therefore, payment flows for "peering" can be confusing. They can represent an outflow (for transit purchases), an inflow (for direct peering with a CAP) or be completely neutral (for settlement free peering with other access providers)

[5] Even if there are multiple peering locations between two networks, the decision to peer rests solely on each of the network's willingness to do so.
[6] Technically speaking, transit clients are also called "peers"

Commercially, peering, either direct or through an IXP competes with transit, and this is good thing. BEREC's analysis of their respective cost profiles is accurate. Direct peering, being a one-to-one relationship offers both parties the ability to use their relative market powers to extract a financial advantage. All access providers, backed by their retail "side-by-side" monopoly positions enjoy a substantial advantage in those negotiations. To prevent these potential abuses of power, BEREC and the European Commission have correctly identified the need for another option to always be available and therefore for peering and transit to be substitutable[7]. A refusal by an eyeball provider to peer with a CAP or even with another eyeball provider (or to subject that peering to uncompetitive financial conditions) can therefore be circumvented through transit purchases.

Anticipating here our response to BEREC's question number 7: "To what extent does the functioning of the peering market hinge on the competitiveness of the transit market?", peering is a one-to-one relationship and Eyeball ISP market powers can fully express themselves in this area, unless transit is an effective "default option". BEREC is absolutely right to underline that "*The peering market is generally taken to function more or less competitively as long as ISPs have a choice of transit providers*."

Tier1 transit providers provide global connectivity (i.e., access to the full routing table) without purchasing transit but through settlement free peering relationships only. Tier2 transit providers rely on a mixture of settlement free peering and transit purchases (from Tier1 operators) to provide the same global connectivity.

3. **Settlement-free peering arrangements:** As we have just seen in the description of transit above, settlement free peering agreements between transit providers are intrinsic to the internet route propagation mechanism. Some (not all) include the use of traffic ratios, with associated payment (to the net receiver of traffic) if the ratio is exceeded. Peering agreements are also fluid. There are numerous peering and de-peering happening across the globe all the time, but they create no issues as long as global connectivity is maintained and the internet is not partitioned. For example, de-peering a network that is also a client of a peer (i.e., his prefixes can be reached through that peer) will not damage the integrity of the global routing table.

   Routing Protocol: Hot potato applies as the default routing protocol to all settlement free peering agreements. "Nearest exit routing" applies to all traffic received from clients and needing to be delivered to a destination that can only be reached through a settlement-free peer. Faced with a choice of multiple peering locations with that peer, the location that is closest to the origin of the traffic (i.e., the client's) will be selected in order to minimize network load.

   Viewed from the other side of the peering relationship, of course, it also means that traffic received from a peer and destined to a client will be accepted anywhere on the network and usually fairly remotely from where the destination is located.

[7] The Polish case underlined this requirement but unfortunately did not verify that it actually was.

"Eyeball heavy" transit providers (i.e., with mainly incoming traffic) carry the traffic most of the way and bear a greater amount of network load. This gave rise to the well publicized "free rider" accusations that CAP's do not participate enough to the financing of networks. We are pleased to see that BEREC did not find any evidence to support those accusations.

However, we now need to describe how transit providers have competitively adapted to manage this "free rider" issue. It will help us understand why integrated incumbents[8], who are trying to validate the use of settlement-free peering between themselves and the rest of the internet are distorting this competitive response. There are three kinds of competitive responses;

a) **Using alternative peers:** Virtually all Eyeball ISP's [9]have multiple transit providers, and their prefixes are therefore advertised and propagated by each of them. The resulting route multiplicity provides alternative peers through which traffic can be sent. Since every network is in direct control of its outgoing traffic, the easiest way to avoid unbalanced traffic with a specific peer is to route that traffic through a different one.

   Integrated incumbents are forcing their eyeball operations to exclusively buy transit internally. Those eyeball networks are therefore "single homed" to their internal transit division, which becomes just as unavoidable as their access division. As a result, there are no alternative peers through which traffic can be routed.

b) **Cold potato routing**: Instead of routing outgoing traffic to the nearest exit peering location, traffic is routed to the further exit point, i.e., to the peering location that is closest to the destination of the packets. The result is a shift of network load from the "eyeball heavy" peer to the "content heavy" transit network. Importantly, cold potato routing does not change the ratio of incoming/outgoing traffic between the two transit networks, but merely re-arranges the locations where traffic is delivered.

   Between transit providers, ratios are part of the broader concept of "bitmile" parity [10] which also includes the locations where traffic is exchanged. As a result, peering policies tend to include a measure of relative network size and reach.

[8] Those incumbents operate both an eyeball ISP and a Transit business, and are claiming to have achieved Tier1 status and provide global connectivity without resorting to transit at all.

[9] Unsurprisingly, the exceptions are the integrated incumbents who force their eyeball businesses to exclusively purchase from their internal transit division.

[10] BEREC refers to those geographical locations (8th bullet point under section 3.2) as needing to be in the same "country". This really needs to be expanded to the "world" as backbone providers' networks span multiple countries and continents. Settlement free peering between them depends on their global reach rather than specific locations within a specific country.

It is worth mentioning here that the application of cold potato routing requires the agreement of both networks. They need to cooperate and exchange MEDS (Multiple Exit Discriminators) to make it possible.

Integrated incumbents are refusing to provide sufficient peering capacity in their home market and cooperate on MEDS. Not only does this lead to QoS degradation for local traffic (increased latency due to tromboning which unfairly advantages incumbents' local transit offers) but it also is a non rational decision, as it would save them costs by shifting network load from their transit network to the originating one.

c) **Client warfare is the most effective response**. It goes without saying that "eyeball heavy" transit providers will attempt to convert the CAPs that are responsible for this heavy inbound traffic. Similarly "content heavy" transit providers, facing a potential payment to their transit peers (due to the provision of ratios in the settlement free peering arrangements) will also try to gain eyeball ISPs clients so as to bypass their peer altogether

This price competition is good for both sets of end users and has contributed to the steady declines witnessed in bandwidth prices.

Integrated incumbents are forcing their eyeball operations to exclusively buy transit internally, whatever the price point offered by competing transit providers.

4. **IP bandwidth is a commodity** –

Route multiplicity directly results from multi-homing and multiple route propagations by transit providers through their clients and settlement free peers. It allows all end clients (CAPs, Hosters or eyeballs ISPs) to make transit purchase decisions independently from each other without losing global connectivity. With global connectivity assured, transit providers effectively become inter-changeable. They all sell the same "global access" product and price becomes the main driver in the buying decision process.

Only with route multiplicity can end users benefit from transit prices that fall in line with the bandwidth marginal cost of production, itself driven by technology advances.

Some transit providers have developed so called "partial transit" offers for regional traffic. That traffic travels less distance and incurs less transport costs. It can therefore be offered at a lower price than "full transit". Even if it can only be offered, by definition, as a complement to "full transit", this is welcome competition and makes perfect sense.

5. **Bill And Keep is essential to route multiplicity**.

Route propagation to settlement-free peers is based on the need to maximize revenue by attracting as much traffic as possible to the receiving client. Peering is a "non transitive" business because there is no economic sense to carry traffic without any compensation. A fortiori, prefixes acquired from peers or clients that would need to be paid to receive incoming traffic would not be propagated to settlement free peers either. This would immediately translate into a net loss situation (receiving traffic for free and paying the destination ISP).

For this simple reason, the introduction of the old CPNP billing method to internet traffic would ring the death bells of the internet. Not only would the revenue loss need to be recuperated from the CAPs and Hosters (huge impact on bandwidth prices and business creations), but route multiplicity would be destroyed to be replaced by route singularity. End clients will no longer be able to make bandwidth purchase decisions independently from each other. CAPs and Hosters will need to become clients of the transit provider selected by the Eyeball ISPs. Price will no longer be the criteria of choice when selecting transit providers and bandwidth will cease to be a commodity.

BAK defines the internet, just as much as the separation between network and service layers.

In summary, a pure transit provider operates in market defined by a set of existing conditions

- Transit networks are avoidable: The system only works because route multiplicity always provides another peer through which traffic can be routed. This is probably the best example of infrastructure based competition cherished by regulators.
- Settlement free peering between transit providers requires similar sizes. The system only works because peer's networks are of similar sizes, and cold potato routing is always an option.
- Clients make independent transit purchase decisions. Route multiplicity provides global connectivity, transforms IP bandwidth into a commodity and allows clients to respond to price competition.
- Transit providers have an economic incentive to balance traffic - As opposed to Eyeball ISPs or Hosters, they *cannot* specialize in either kind of end-clients. The economics of the transit market forces all transit providers to balance their client mix properly and maximize "on-net" traffic (i.e., traffic between their own clients, i.e., CAP or Hosters towards Eyeball ISPs) for which they collect revenue on both sides and minimize the "off-net" traffic (peers to/from clients), for which they collect revenue on one side only.

There are **three essential conditions** for transit to be an effective default option: Connections must be **congestion free** and **local** and **end users' prefixes propagation must not be artificially withheld**.

Question 5 (Chapter 2): Do you agree with the classifications of CDNs as outlined above?

BEREC's description of CDN is accurate. Cogent will therefore only make five observations.

1) In point 2.4.d), BEREC indicates that revenue is usually billed on per Mbps basis (i.e., capacity per second). To Cogent's knowledge, we believe however, that revenues are mostly billed on a per Megabytes basis (i.e., total volume of data transmitted).

2) To our knowledge, there is only one CAP that is big enough to deploy its own CDN, i.e., Google. But even Google has to buy a certain amount of transit in order to deliver their content to the access networks they do not reach directly.

3) The CDNs business model is dependent on the transit market. CDNs need to buy transit for two reasons: They need to provide global distribution from their CAP clients to all eyeball ISPs, including those that do not connect with them directly and they need to regularly update their cache servers. The transit market is the canvas that enables the CDNs business model.

4) CDNs do not provide global connectivity to the eyeball ISPs but only allows them to access the content from by their CAP clients. From the eyeball ISP standpoint, CDN is just like direct peering: it is a "one-to-one" relationship that only delivers route singularity, in which Eyeball ISPs market powers can fully be expressed.

5) CDNs compete with transit and direct peering as a way to distribute content to eyeball ISPs. Each method has its own cost structure, commercials, technical advantages and quality mix. CAPs must be able to choose either of those distribution methods on their own merits. They should not be force to pick one over the other because the price/quality of either is manipulated by Eyeball ISPs.

To ensure a level playing field, transit, as the sole provider of global connectivity, needs to always be an option to reach CAUs. Transit needs to be substitutable to CDNs or direct peering.

Question 6: To what extent are requirements regarding traffic ratios still important in free peering arrangements?

We discussed this in our responses to question 1 and 4 above. To summarize our response:

1) Due to the causality of traffic, ratios have no place in relationships between networks of different kinds. They would in fact constitute a violation of net neutrality as they are used to limit CAUs ability to access all applications, contents and services of their choice.

2) Between networks of similar kinds, ratios are only one side of the argument, with the other being the locations of the interconnections. Their relevance is therefore subject to a broader definition of "bitmile parity".

Question 7: To what extent does the functioning of the peering market hinge on the competitiveness of the transit market?

We discussed this above in our answers to question 4. To summarize our response:

The proper functioning of the direct peering market is entirely dependent on the substitutability of transit. Direct peering is a one-to-one relationship and the only competitive pressure that can be applied to the potential abuse of market power from eyeball ISPs is the availability of a transit option.

This is precisely why integrated incumbents are preventing their eyeball subsidiaries to behave independently and purchase transit away from their internal transit operations. At the same time, the transit division is limiting the incoming traffic under the cover of a ratio driven peering policy. The combination of these two results in the ability to extract uncompetitive prices from CAPs for direct peering interconnections.

This is precisely why regulators should always ensure that transit is always an option.

Question 8: Does an imbalance of traffic flows justify paid peering?

We have discussed this already in our responses to question 1. Traffic flows between eyeball ISPs and the internet are structurally unbalanced.

Regulators should focus on ensuring that transit is always an option so that a proper state of competition can exist between the various methods of delivering the traffic requested by the end users. Paid peering is only but one of those methods (together with transit and CDNs) and the market will dictate if payment is needed as well as the level of that potential payment.

Question 9: Does paid peering increase (number of contracts and volume handled under such contracts)?

Question 10: To what extent does regional peering increase in relevance and affect transit services?

The overwhelming majority of Cogent's peering relationships are through dedicated connections and we therefore cannot comment on the relative position or market share of those regional peering exchanges. Regional peering may or may not increase, but what is constant is that they provide normal competition to the transit market. This competition is healthy and contributes to price declines in the transit and in the peering exchange markets; it increases overall benefits to end users.

Question 11: Are any important services missing from the list of services provided by IXPs?

No comments

Question 12: Are there any further developments regarding IXPs to be considered?

No comments

Question 13: Should in future Europe evolve to have more decentralised IXps closer to CAUs?

No comments

Question 14: Will traffic classes ever become available in practice on a wide scale?

Question 15: Will interconnection for specialized services be provided across networks?

As mentioned above, every network only has control of its outgoing traffic. QoS can be implemented in that direction. However, networks have no control on their incoming traffic. QoS can therefore not be dictated on the traffic received from the internet. The emergence of QoS and traffic classes across networks will therefore only be possible if this is forced on all CAPs, either externally or through generalized network congestion.

Traffic classes are responses to congestion. There is no congestion on the Internet, except at the gates of Integrated Incumbents, who are setting those up deliberately to create artificial scarcity. Increasing network capacity is cheaper than establishing traffic classes.

The market will dictate whether traffic classes eventually take off or not. But that choice needs to be made on a level playing field. Specifically, we believe BEREC should focus on ensuring that the interconnection market is competitive and is not distorted artificially by some players whose ultimate goals are to transform the best effort internet into a dirty road.

Question 16: Will other solutions for improving QoE like CDNs become more successful rather than traffic classes?

No comments

Question 17: Which of the factors impacting on the regionalization of traffic is most important: language, CDNs, direct peering?

No comments

Question 18: Are any further issues missing?

No comments

Question 19: Given the cost reductions and the economies of scale and scope observable in practice, why do network operators call for compensation?

BEREC correctly identifies that technological advances and price declines in the router and optical markets have allowed transit prices to decline accordingly. While this is probably outside the scope of BEREC interconnection commentaries, Cogent would only say that those prices could have declined even faster if the router and optical market were not in a situation of virtual duopoly.

The internet is the telephone network's worst nightmare. They can no longer leverage their network monopoly into application delivery. The plethora of services available on the internet is replacing their legacy revenues. Their survival hinges on their ability to innovate, replace those revenues with services and applications and integrate vertically. Unfortunately, their innovation capabilities are low.

The call for compensation, justified officially by the need to invest in the network is a smoke screen. Payments would be minimal compared with the revenue that needs replacing and would go unnoticed in their P&L accounts.

However, the payment's secondary effects are the important ones. They will act as a barrier to entry for services and applications created on the internet ("*off-net services*") and will increase their costs so substantially that it would actually reduce innovation. On the other side, it will allow eyeball ISPs to develop their own internal services ("*on-net services*") shielded from this now reduced competition. Integrated incumbents can then integrate vertically (including by acquiring "*off-net*" competitors are rock bottom prices), behind the walls they have successfully built around their CAUs.

The call for compensation is hiding the much larger potential revenues that eyeball ISPs could generate from discriminating against "*off-net services*" and selling their own "*on-net services"* to their own CAUs.

Cogent has no issues with Eyeball ISPs diversifying into application and services creation. We believe however that allowing them to do within a walled garden and without being subject to proper competition from "*off-net services"* is a matter for regulators to take very seriously.

Question 20: Do you subscribe to the view that CDNs lead to improvement of QoS without violating the best effort principle?

CDN is only but one method of delivering traffic requested by end-users. It needs to operate on a level playing field with other methods of distribution. CAPs and Eyeball ISPs need to choose between those methods on their own merits.

Question 21: Is there a trend for CDNs to provide their own networks (i.e. integrating backwards)?

No comments

Question 22: Is there a general tendency for eyeball (CAU) ISPs to deploy their own transit capacities and long distance networks or even to become Tier-1 backbones?

Only a handful of European incumbent eyeballs ISPs (DTAG, FT, Telefonica, Telecom Italia and Telia) have historically deployed long distance networks. They are all using them principally to transport traffic requested by their CAUs, in the various countries where they operate. Only one (Telia), to Cogent's knowledge has also developed neutral competitive transit activities and stands as a Tier 1 provider on its own merits.

DTAG and FT have become Tier-1 in 2010. However those are fake Tier-1's. The size and reach of their long distance networks as well as their level of connectedness (number of independent connected ASs) would never have allowed them to become Tier 1 on their own merits. They both have achieved this status by unduly leveraging their SMP in the retail market (forcing their eyeball ISPs activities to exclusively buy transit internally, thereby making their transit operations just as unavoidable as their access one) into a preferential position in the transit market.

This trend is likely to continue unless regulators ensure that transit is always substitutable with peering at the gates of the eyeball network.

As a reminder, it takes two to peer and it is not unconceivable – if the peering relationship is unsatisfactory - that real Tier-1's depeer the fake ones. The internet would then be partitioned and it will be left to regulators to decide how best to restore global connectivity for the impacted CAUs.

Question 23: If an eyeball ISP becomes Tier-1 provider, does this increase the eyeball's market power on the interconnection market because there are no alternative Tier-1 providers to reach the customers of this eyeball ISP?

Absolutely yes –

In fact, the way BEREC phrased the question already reveals the confusion about the two activities. Access and Transit are two very different economic sectors in the internet ecosystem. Incumbent operators can of course develop transit operations and try to become Tier-1's on their own merits. But the preservation of a competitive level playing field in the transit and access sectors makes it imperative that each operation behaves independently. Structural separation of those activities makes sense.

Specifically, the exclusive internal purchase of transit transfers SMP from the access to the transit division and reduces route multiplicity to route singularity on purpose. Transit is no longer an option to reach the customers of this eyeball ISP. The competitive landscape is altered substantially.

The next step is of course for the fake Tier-1s to demand payment from their peers under the application of traffic ratio imbalances in the settlement free peering agreements. Of course, those ratios are structurally unbalanced (they reflect CAUs sum total of surfing behaviors), the transit division is unavoidable (to the opposite of a normal transit operator) and there are no alternative routes to reach the eyeball ISP (because of the exclusivity granted to the transit division).

The fake Tier-1s have already succeeded in establishing a walled garden around their CAUs. CAPs and proper transit providers' only choices are to either suffer QoS degradation due to congestions on the peering interconnections to the fake Tier-1, or to pay the integrated incumbent for connectivity (through direct paid peering contracts or because of ratio overage) at non competitive prices.

Question 24: Will Art. 5 become more relevant as some large Eyeballs have equally qualified as Tier 1 provider not having to rely on transit anymore?

Art.5 is certainly the core of the issue. However, there is another logical solution that stems from following the logic of the internet to its natural conclusion. We will review both below

Art.5 was written in pre-internet times to validate Incumbents' right to receive call termination payments under the old CPNP billing system. Regulators have since fought a constant battle to bring those payments down in line with incremental costs rather than arbitrary levels.

The Internet is different. As we have seen, traffic is always caused by CAUs and CAPs only respond to download requests. CAU's internet subscription covers both outgoing and incoming traffic. As we have seen above as well, the Internet's very existence (the prefix propagation system) rests entirely on the BAK billing model.

NRAs have been shy to use Art.5 in the various disputes between Cogent and DTAG or FT[11], because of its implicit right to receive payment for connectivity. Cogent engaged the NRAs because the only path available to deliver CAUs requested traffic was (and still is) via the deliberately congested peering interconnections with those fake tier-1s. In other words, the "default" transit route was not an option because the incumbents 'eyeball ISP divisions were (and still are) exclusively buying transit internally.

BEREC appears now to properly place Art.5 AD in the broader context of Art.8 FD and Rec. 8 AD:

"*However, under certain circumstances conflicts may arise when one party denies a plea for interconnection and thus would be able to take customers hostage…In these cases NRAs may have to take action in order to promote and defend fair competition, investment, innovation and consumer welfare in the sense of Art 8 FD and may decide to impose obligations to interconnect."*

BEREC recognizes that Art.5 cannot be used to force eyeball ISPs to use any particular type of interconnection (transit or peering). However it is clear that if one specific type is denied (peering), then the other (transit) must be made available. Since transit becomes available if either party buys it, the question becomes: who should be obliged to do so?

Art.5 clearly states that this obligation can only be imposed on "*undertakings that control access to end-users*". Transit ISPs do not control end users. The obligation to buy transit and restore end-to-end connectivity rests squarely on the eyeball ISP's shoulders.

It follows naturally that if a real Tier-1 provider refuses to peer with a fake one, then regulators should use Art.5 to ensure that the eyeball ISP restores global connectivity by buying transit from another independent transit provider.

[11] BNetzA acted effectively as an arbitrator and, as relationships improved with DTAG, did not have to settle the point of principle. Arcep avoided the issue altogether and encouraged Cogent to turn to the competition authority instead.

There are three essential conditions for transit to be an effective default option. It needs to be congestion free, local and the propagation of end users prefixes must not be artificially withheld.

The separation of the local access network from the access to the global internet. This would allow end user to have the choice of multiple IP upstream providers running on top of their local access network subscription. Consumers can then enjoy the benefits of open and well informed high quality offers, at little to no switching cost**s.** Just as the internet content creation has been permitted by the separation of network and applications layers, its logic should be applied to the last mile access providers. Much work has already been done, specifically in Working Group 3 of Commissioner Kroes' internet industry consultation in 2011, to show that separation was possible in the retail sector.

Access networks should be neutral and open to all forms of internet connectivity offers. In this business model, the last mile provider should offer network connectivity only at the hardware level and users' subscription cover only the costs of providing the last mile network access. This would be a regulated monopoly under NRAs' control, with no permission to provide internet access or any other additional services.

Numerous participants have testified that such access network monopolies will enjoy easy access to financial markets and that the roll-out of the European Digital Agenda will be greatly enhanced as a result.

Independently of their hardware connections, users will then be free to choose Internet providers directly. Alternative providers can all compete for their business as switching becomes extremely easy. Only then can Regulators be assured that users will benefit from bandwidth's rapid price declines and from the absence of conflict of interest generated by "the need to finance the network". Incumbents will face a clear choice of either becoming a true application, content or service competitor or a basic, entirely neutral network operator. While this solution might sound revolutionary, it actually is not. Other regulators have already travelled this route, by separating creation from distribution in the areas of water, energy or electricity.