Dear authors,

I write this e-mail in a personal capacity. It represents my own opinion and not that of anyone else. I write it because of my interest in IP interconnection for the last 2 decades. An explanation of peering and transit. I wrote 16 years ago is at https://arstechnica.com/features/2008/09/peering-and-transit/. I also initiated the BEREC-OECD meetings on IP-interconnection in 2011 when I was employed at the OECD. I realised then that regulators, lawyers and economists had good knowledge of telephony interconnection and the dangers of termination monopolies, but little knowledge of how the Internet had a different model. Those meetings brought the peering community in contact with regulators. Then telecom operators tried to hijack the WCIT discussions to create a mandatory termination monoply for them. It made clear that much of what incumbent telecom operators and their economists had written on the subject was incorrect. Unfortunately the same myths that were brought forward by telecom operators then are today again expressed in the debates on "fair share", "network fees" etc. The academic research that is supposed to support these opinions is equally flawed as it was in 2011.

Thank you for you work on IP-interconnection. The report is well written and I support its overall conclusions. There are some elements that could have some review, which I address below.

- The report makes a distinction between IAS, CAP, CDN etc. It would be good to make clear that although there may be differences in business model, the system of IP interconnection doesn't see any difference between one ASN and another. The business model is irrelevant to how traffic is routed. The investments made by various ASNs also reflect this. There is no IAS who needs global connectivity if it only serves consumers in one country. A CDN or gaming firm will need to invest globally, because their customers are around the world

- The Term Tier 1 providers is a misnomer that just doesn't want to go away. A network that doesn't pay any other network for transit sounds cool, but it also means that if something fails that network needs to have enough capacity and locations to handle the problems. A network that buys transit from multiple transit providers can depend on its transit providers for alternative routes. As I wrote in 2008:

> Tough at the top: word about Tier 1 networks
>
> Tier 1 networks are those networks that don't pay any other network for transit yet still can reach all networks connected to the internet. There are about seven such networks in the world. Being a Tier 1 is considered very "cool," but it is an unenviable position. A Tier 1 is constantly faced with customers trying to bypass it, and this is a threat to its business. On top of the threat from customers, a Tier 1 also faces the danger of being de-peered by other Tier 1s. This de-peering happens when one Tier 1 network thinks that the other Tier 1 is not sufficiently important to be considered an equal. The bigger Tier 1 will then try to get a transit deal or paid peering deal with the smaller Tier 1, and if the smaller one accepts, then it is acknowledging that it is not really a Tier 1. But if the smaller Tier 1 calls the bigger Tier 1's bluff and actually does get de-peered, some of the customers of either network can't reach each other.
>
> If a network has end-users (consumers or businesses), it's probably in a better business position than a Tier 1 or a pure-play transit provider, since having end-users provides stability to a business. Autonomous Systems can switch within seconds, but end-users are stickier customers. Churn is less of a problem and revenues are therefore more stable and easier to base decisions

on, since prices don't have to drop on a monthly basis. So an end-user business, combined with a bit of transit is, therefore, ideal for a network provider.

With regards to Tier 1 status, I also want to mention that a large telecom operator who is an ISP to millions of consumers, but also claims that it is a Tier 1 is a dangerous situation from a resillience point of view. It mean that if it doesn't have enough capacity it can't fall back on other networks to alleviate its issues. We saw this in the Netherlands when Deutsche Telekom misconfigured the routing of its Dutch ISP T-Mobile and all traffic went through Frankfurt, where Deutsche Telekom apparently lacked capacity to reach Dutch government websites and universities. https://rudolfvanderberg.medium.com/t-mobile-nl-routed-all-internet-traffic-through-germany-and-broke-the-internet-for-small-firms-a176855d2b0

I would warn against using Sandvine as a source. It has limited insight into EU ISPs. In addition it doesn't matter what is in an IP-packet. An ISP should handle it regardless. Total traffic is also irrelevant, because networks are dimensioned for peak capacity. It is analogous to road traffic, where the problem is with rush hour and not how many kilometers a taxi drives. The nice thing of traffic on fixed and mobile networks is that with 3-7Mbps peak per subscriber there isn't an ISP in Europe who should have difficulty with handling that traffic. The 50Tbps peak traffic that ARCEP reports in France can easily be handled by 200Tbps routers of Nokia.

Inbound and outbound traffic ratios are quite meaningless too. For the same peak a 1:1 or a 5:1 ratio requires the same network equipment (some caveats apply)

Your description of on-net CDNs might benefit from a mention that an on-net CDN is an extension of peering. Why pay a third party for transit, or carry all the traffic from Paris, Frankfurt or Amsterdam when it can also be on a server close to the consumer. It saves the ISP in backhaul capacity upgrades and makes the service more resilient.

With regards to paid peering, there should be a mention that some incumbents claim to sell transit to other networks, which is a just a termination payment disguised as transit. As mentioned in this blogpost https://rudolfvanderberg.medium.com/request-for-retraction-of-netneutrality-and-high-speed-broadband-networks-evidence-from-oecd-d33ce2f8b749

Addendum 1: An employee of a European content network send a message that the tariff charged by Deutsche Telekom for terminating traffic on its network is a blended price of transit and termination payment. So it may look 20cts per megabit per month capacity, but it really is a termination fee that is high and an additional lower tarriff for transit, which the other may not want or need.

Any network that has significant traffic for Deutsche Telekom finds that commercial transit networks towards DT are saturated (an experience that Deutsche ForschungsNetz had too). The network than requests Deutsche Telekom for peering. Deutsche Telekom Wholesale will only offer transit. The price of transit is different for different AS-numbers, based on whether they are a telco like Deutsche Telekom, CDN or hosting provider. This is odd, because networking protocols can't see a difference and there is no cost difference. Hosting apparently pays more than CDN and telco. The network than tells Deutsche Telekom it has enough transit capacity and doesn't need more.

Deutsche Telekom however insists, it's either transit or nothing. The transit price for 100Mbps may be 20 euro per month. The hosting network only has 25Mbps of traffic in either direction with DT. The price for terminating traffic translates to around €0.80/Mbps/per month. 4 times more

than I wrote based on WIK. The content provider can then choose to use the transit or not. If the content provider wants only 50Mbps than DT will tell it the price is 40cts/mbps so €20/mbps/month too.

However requesting to only pay for the traffic that terminates on DT, paying only the blended price of €0.20/mbps/month is met with disbelief. It still costs €20/month for 25Mbps. If the content provider doesn't agree, the only alternative is to route through networks DT has a friendly relationship with and that have a similar price point. This is similar to the Swisscom-Init 7 case referenced below. Mind you, Swisscom lost in court for exacting a termination monopoly with a revenue split with Deutsche Telekom for all traffic destined to Swisscom routed over Deutsche Telekom.

- The issues in chapter 6 are not so much issues in IP-interconnection. They are reflections of a termination monopoly being abused. It is not a problem that an incumbent ISP doesn't want to peer with another network. It is a problem that it refuses to purchase enough transit to handle incoming traffic from networks it doesn't want to peer with. The competition problem is that consumers can't see what causes the failure of the service they try to use and blame it on the website, game or video they try to use, instead of on their ISP. Changing ISP is also harder than going to another website and as a result the ISP doesn't feel competitive pressure not to exploit his termination monopoly. The VodafoneZiggo merger case had a requirement for unimpeded transit through 2 transit providers. Such a remedy could be mentioned here. The cases in the annex give an interesting insight. I would suggest that the way DT forces DFN to pay a termination fee, because otherwise German Students couldn't follow online classes during Covid lockdowns, deserves a mention too.

I'm happy that the report doesn't mention two-sided markets and the mistaken industrial organisation economic literature on the subject. An IAS is not a two-sided market and IP interconnection can't be described as a two-sided market. Yes, there are some famous economists who have tried to fit a two-sided model on IP-interconnection, but that doesn't make it correct, nor is there any empirical support for the predictions of those theories. After over 20 years of theoretical literature not fitting with reality, it is fair to say the theory doesn't match the empirical data and should be thrown out.

I wish you all the best in reviewing all the comments and hope the final report will reflect the current version, with some improvements.

Sincerely,

Rudolf van der Berg

Netherlands.