Net neutrality observed in more detail: Influences on *end user experienced neutrality* of Internet based services

Author: *Sietse van der Gaast,*
Organisation: *Stratix Consulting,*
Email: *sietse.vandergaast@stratix.nl*

Organisation address: Villa Hestia - Utrechtseweg 29, 1213 TK Hilversum, the Netherlands
Telephone: *+31 6 27106263*
Fax: *+31 35 6222029*

Keywords: net neutrality, network, Internet, business model, regulation

Abstract

This article categorises influences on end-user experienced neutrality of the Internet by surveying the influences on end-to-end Internet services, not limiting the scope to specific physical network boundaries or functional layers.

It elaborates on the theory of net neutrality, and the difficulty of a clear definition due to varying end-user roles, and implicit definitions of concepts such as 'Internet' and 'network access'. The article also describes influences weakening net neutrality in practice, the responsible stakeholders and the economic side of such influences.

The playing field is modelled by breaking up the potential influences by different functions that provide end-to-end communication on three different axes: their position in the network topology, in functional layers and in network geography. Examples are given of using this model to map and investigate net neutrality threats.

Finally an alternative definition of net neutrality is given, and recommendations are given to avoid pitfalls in regulation.

1. Introduction: discussing the pursuit and the spirit of net neutrality

Strong sentiments and bold statements dominate the net neutrality discussion. Those in favour of both freedom of speech and the free market are having a hard time, as these principles seem to pull in opposite directions in the case of net neutrality.

This article attempts to better describe the notion that is indicated by the vague term “net neutrality”. It tries to untangle the different issues, and to uncover the real challenges that the Internet of the future faces. An attempt is made to describe end-to-end experience of a service, to avoid disregarding important influences in end user experienced neutrality of the internet. Therefore, the scope of this article is not limited to specific physical network boundaries or functional layers.

1.1 Background of the current net neutrality discussion

Companies offering “common carrier” services such as the transportation of mail, speech or data are commonly expected not to tamper with the content of the information that they transport. Neither are they expected to limit the freedom of choice of the customers they serve with regard to their selection of the recipient of this content. For speech services, this principle has been given a specific legal basis, but for Internet services the legal aspects of this matter are still a greenfield. Intuitively, people are expecting ‘the Internet’ to be open and neutral, and at least until recently the different parties that transport data from end user to end user seemed to honour this principle, albeit on a best effort basis. The history of this expectation lies in the end-to-end principle of the original Internet architecture, where communications protocol operations are primarily defined at the end points, or as close as possible to the resource being controlled. The result is that ‘intelligence’ is located mainly at the edges of the network rather than in the core. This architectural feature is considered by many to have been a key driver of the growth and success of today’s Internet, and to have facilitated an open environment conductive to the spectacular levels of innovation seen in online applications, content and services [17].

That the theoretical openness and neutrality is not entirely met in practice, mostly for simple technical reasons, is widely accepted as long as they are not deliberately countered. Providers treat all email equally, but in practice it is very well possible that an email sent to an address that is hosted by the same provider as the recipient’s address is delivered a faster than an

email sent to an email address that is hosted by another provider. However no one expects that an email may be blocked, deliberately delayed, or its content altered, just because it is sent to an email address hosted by a competing provider.

The increasing diversity of the services offered via Internet, the popularity of broadband Internet services, and the bundling of services by service providers, all create an incentive for ISPs to interpret the 'best effort' principle in a looser manner, or to even abandon the principle altogether. Consumers (and consumer rights groups on their behalf) are noticing that some ISPs are delaying or even blocking certain categories of data traffic.

This type of behaviour leads to questions about whether service providers should have the freedom to make their own decisions regarding traffic regulation [17]. Even if we suppose that the competitive pressure is sufficient to discourage unacceptable behaviour, it still is a valid question whether rules are needed to safeguard transparency and thereby enable proper competition in the marketplace.

1.2 Net Neutrality in theory

The open Internet has proven to be an excellent platform for initiating and deploying new innovative services, and also for new ways to offer already existing services. It is not always clear what exactly the properties are that makes the ideal net 'open' or 'neutral'. This also depends on the perspective of the user. In literature there are many different definitions of the term 'net neutrality', most of which contain in some form the following three propositions::

- *Content*:
  As with telephony, the access providers do not have any influence on the informational content of conversations, and the choice of persons between whom the conversations are made; for Internet this also does not seem desirable.

- *Connectivity*:
  Every (public) Internet end point should be reachable from any other (public) Internet end point.

- *Quality*:

An end user should be able to enjoy a service offered on Internet with a comparable perceived quality as another end user that has similar Internet access properties and similar end user equipment.

Discussions around net neutrality are commonly based on a very simplified picture of the Internet. A development such as the blocking or delaying of specific types of Internet traffic is being compared with the utopian Internet where every service is available for all persons, with the same quality, as in figure 1.

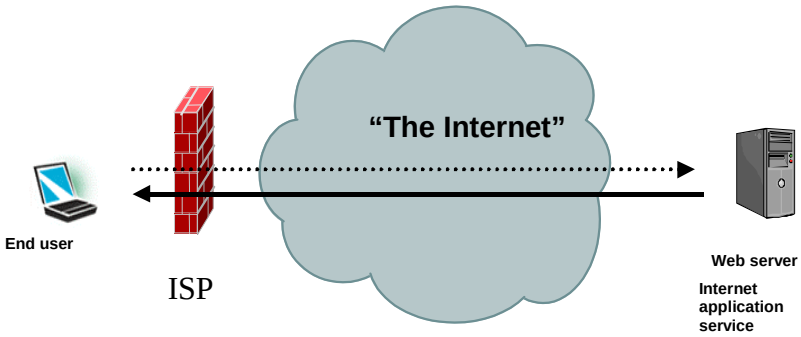

Figure 1: Internet services in theory as perceived by the user: The ISP as gateway to a world wide network.

2 Net neutrality in practice

The way the Internet works is in practice - not surprisingly - far more complex then the idealised picture of a simple network as described in the introduction. This chapter describes the complexity by elaborating on the role of user, the different aspects that influence the behaviour of Internet services, and the problems related with the definition of Internet access.

2.1 Different faces of the end user

The roles and the interests of the end user are often ambiguous and raise different kinds of questions related to the notion of net neutrality. We can distinguish the following roles:

- *The end user as a consumer* [6].

  The user that sees Internet Access as a product. Ideally it should be possible to choose between more than one Internet access subscription offered in a fair, open and free

market, and to compare subscriptions in a not too complicated way on aspects like capabilities, quality and price.

- *The end user as a citizen* [7].
  This relates to discussion on citizen rights. Since communication, not only with companies but also with the government, has become almost impossible without Internet access, one can ask the question whether open Internet access should be a right [8][9], and if so, how we can ensure this right. Another issue that relates to this role is the question whether the application of Internet filters can be in conflict with the right of freedom of speech and freedom of information retrieval [11].

- *The end user as a content producer or service provider.*
  One of the success factors of the Internet is that it allows for new business ideas to develop from very small implementations and deployments. Is the growing importance of the 'long tail'[10] in the content- and services offerings making a new vision on regulations necessary, when compared to the more limited conventional market place?

- *The end user as a network administrator.*
  The end user may operate a local LAN and therefore act as a small network operator. This relates to issues like: Can the end user as a network administrator, connect all kinds of equipment, including servers? Can the end user share his access to the Internet with third parties (for instance via open WLAN)?

The relevance and importance of the different issues and questions in relation to the concept of net neutrality depends heavily on which end-user roles are considered. To avoid solving one problem by creating another, it is important not to focus on one end-user role, but to take all of these roles into account.

2.2 Different faces of the network

Data exchange over the Internet is predominantly characterised by a so-called best effort approach, in which we assume every link in the chain of network elements to handle the Internet traffic as efficiently as possible, but without any form of guarantee. Dependent on the structure of the networks and network elements, this by itself results in a different user experience for users in different networks on different times of day.

Figure 2 shows that one relatively simple service such as viewing a single web page relies on the quality of many connections with many servers, via one or more chains of networks, and via a number of gateways and firewalls, and therefore bottlenecks.

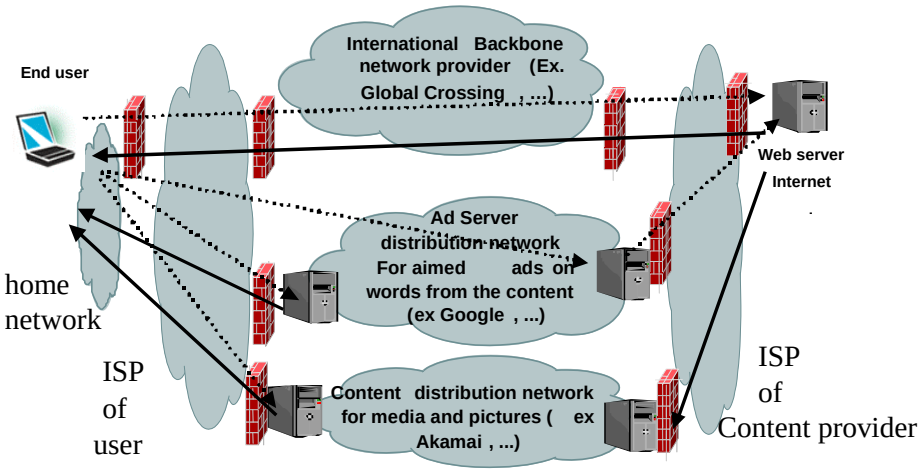

Figure 2: Internet services in reality: many interconnected networks, servers and gateways

Aspects influencing the end user perceived quality of an Internet service could be divided in the following categories:

- *Influences in the end user domain*

  Due to the large variety of user equipment, requesting the exact same Internet page on two different devices can lead to two totally different quality perceptions due to differences in hardware, software, or (home) network: 'net neutrality' does not automatically equal 'web neutrality':

  - Hardware aspects like screen size, speed of the video card used, sound features, memory and speed of the CPU, etc. can influence user experience.
  - Software aspects can also have great impact. Different browsers are known to display identical HTML-coded pages in a different way. The use of spam filters, virus scanners or running other software simultaneously on the same machine can delay web services.
  - The connection in the end user domain is important: it makes a difference whether the end user is connected via an Ethernet cable (or even glass fibre), or wireless via a WiFi network or Bluetooth. It also matters whether other people are sharing the same (home) network and transport their traffic simultaneously over the same connections.

Similar influences of course also apply to other Internet services than web browsing.

- *Influences in the ISP domain*

  The (maximum) bandwidth that the ISP offers is only one of the factors that determine the data transmission speed from and to the end user. Besides more obvious technical factors related to the type of access network used (like: mobile networks are more susceptible for bandwidth variations), the actual network quality (packet loss, delay and response time) and the number of users actually using the network (overbooking) are also of importance. This is especially true in the case of access networks that use an overbooked 'shared medium': in this case the same cable or the same frequency space has to be shared by multiple users. If an ISP is not able to segment its network to accommodate all network users and usage, this may lead to a very different perception of the same Internet access subscription, depending on time and place. Another aspect that is becoming more and more important is the symmetry of the offered connection. Especially in situations where a 'shared medium' is used, ISPs tend to offer high downstream capabilities at the expense of the upstream capabilities. Additionally, the ISPs are offering (mandatory or optional) services that can influence the data transmission speeds, like spam or virus filtering, local caching, and prioritising certain types of traffic.

- *Influences in the backbone network*

  Depending on where in the world wide network topology the service is offered, the traffic has to take a longer or shorter route to the end user. International traffic to and from Internet Service Providers (ISPs) is routed via one or more backbone networks that may be owned by other parties. Both ISPs as well as backbone network providers sometimes limit or block certain routes temporarily, due to illegal or destructive amounts of traffic offered by the peer network(s) or illegal or destructive content detected. The – at first sight very valid – reason for this is to avoid (possible future) network congestion. However it is not always straightforward what exactly can be classified as 'illegal' or 'destructive'.

- *Influences in the domain of the party offering the content or service on the Internet*

  The perception of the quality of the service offered over the Internet also depends strongly on how the servers hosting the Internet pages or Internet services are dimensioned, and how they are connected to the Internet.

- Some services are unavailable for some users, for instance because the service is only licensed for use in a certain geographic area. Viewers in the Netherlands can not, for example, watch Top Gear episodes using the BBC Iplayer. A relatively recent phenomenon is the blocking of content for a certain geographical area purely because the cost of outgoing bandwidth for serving users in that area is larger than the benefits from targeted advertising. After all, only a part of the world-wide Internet users are interesting for advertisers [15].
- It is not uncommon that a service or a web page is built out of segments that are provided by different parties. This is not always noticeable to the end user, but may result in a different perceived quality. An example is the distributed caching of content by companies like Akamai as also described in Figure 2.

Also some large content or Internet application providers have their own world-wide backbone network, with more or less specialised servers near different end user groups. Examples are Google, Facebook and Microsoft with servers and databases world-wide, but also providers of online games like Blizzard have gaming servers in various continents to enhance gaming experience of the end users.

3 Broader technical and economical perspective on Internet Access

An ISP offers end-users a data connection with the capability to access other devices connected to the Internet in a product generally known as “Internet Access”. The end-user may actually perceive this as a product that should provide global access to content and services offered on-line.

This chapter describes why it is important to take a broader scope when considering network neutrality issues, by zooming in on the concept “Internet Access” itself and its relation to other services, and some considerations regarding traffic categorisation in types of content or types of service. Finally, costs and benefits of traffic discrimination are given for major stakeholders.

3.1 Internet Access and its relation to other services

Net neutrality discussions are mostly focused on the (IP) connection offered to end-users as ‘Internet access’ service. This limitation seems sensible, but there are several good reasons not to limit the discussion to this specific connection but to put the issue into a broader scope:

- *Influences outside the "Internet"*

  As mentioned earlier, influences on net neutrality also include the user equipment and the home network.

- *Direction and symmetry of access*

  Access not only means access *from* the end user but also access *to* the end user. This implies that it is important to also take into account aspects like symmetry, upstream capabilities and the ability to connect to a specific end user (for instance the choice to choose a fixed public IP address or DNS entry).

- *Relation to other services that use other access networks*

  Many services are delivered over an IP connection to a home, but not necessarily over the Internet part of that connection. This is for instance the case for many Voice over IP and IPTV services. The fact that similar services can also be offered over the Internet makes it very tempting for a provider to actively influence the Internet traffic related to similar services [1], in order to ensure the superiority of the quality of his own separately offered (revenue generating) service. This may even be the case if the provider offers non-IP based services such as analogue cable TV.

- *Relation to other services that use the same access network*

  Another dependency is the possibility to share the allocated bandwidth for Internet-access with other offered services such as an IPTV subscription. This could imply, for instance, that the maximum bandwidth offered for the Internet connection can only be achieved in reality when no one in your household is watching (IP) television. This is visualised in Figure 3.

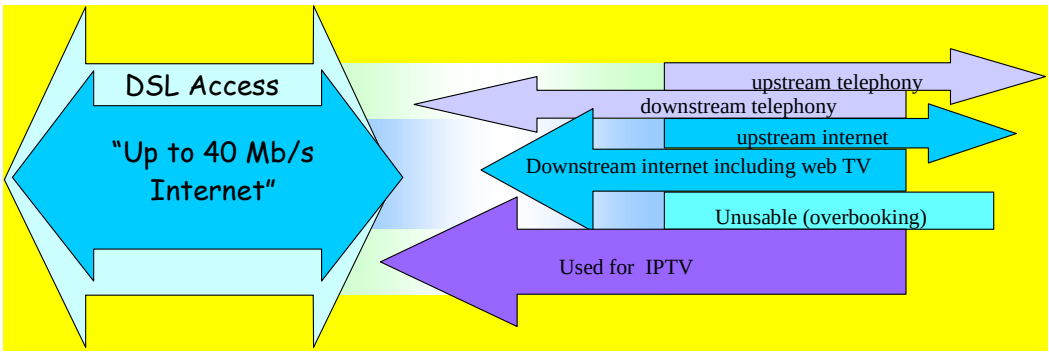

Figure 3: Different services using the same 'access pipe' that is administered by one of the service providers: not everything is Internet, but may be related to the handling of Internet traffic.

- *Minimum limitation of the "Internet" unclear*

  The boundary of the "Internet" in the net neutrality propositions described in 1.2. is not always clear. What does "access to the Internet" at least consist of? Is it the access to *all possible content* on the public Internet? Is it the access to *all possible services* (including for instance voice and streaming video) that can be offered over the public Internet? What are the limits of what may be considered *'public'* or *'publicly accessible'*?

- *Behaviour of "Internet access"*

  Does Internet access also imply guaranteed behaviour of incoming traffic and mandated behaviour of outgoing traffic according to standard Internet protocols?

- *Minimum access capabilities unclear*

  The word 'access' is complicated and multi-interpretable in itself. Many services are already acceptably accessible (to a degree) over a connection with very little bandwidth and very long response times. But some services are only usable over a connection with a large bandwidth, some only need a fast response time, and many need both [5]. For most currently used services the downstream direction (to the end user) is the most important, but for services like video conferencing or hosting web pages the capabilities of the upstream direction are equally or more important.

- *Neutrality does not stop at the access networks*

  Even when a service is clearly identified, the way it is perceived can vary per user, even if they use similar equipment and access capability [4]. This is due to the fact that one service is not always related to a single service provider. As explained above, what

appears to be one Internet service for an end user, is in many cases in fact a set of different services, offered in different domains by different parties in a complex system, with possibly indirect relations, unclear responsibilities and no clear single point of contact. For instance, a web page viewed by user A could be much slower in responsiveness and longer in download time than the same web page viewed by user B, caused by difference in the targeted ads included in the web page. This may have several causes. It is possible that the page viewed by user A contains ads with animations which consume a lot of (data) processing power compared to simple ads in the same page for user B, causing it to appear slower. But it is also possible that the ISP of user A has (for technical or economical reasons) a slower connection with one or more of the used services or content providers than the user B's ISP. Therefore it is important to also consider peering agreements of ISPs with peer networks and of ISPs with specific providers of content (including content caching) and applications on the Internet (such as Google, Akamai and Microsoft). Apart from that, the behaviour of (large) Internet content and application providers with regard to the 'neutral' provision of their services and content to different end users should also not be neglected.

3.2 Distinguishing service types and content types

The core of the net neutrality discussion is formed by the question whether or not to promote or demote certain types of data traffic over other data traffic, based on the type of service or type of content, especially in cases where the network cannot be sufficiently overdimensioned. The claimed advantage of promoting and demoting certain data traffic is to better guarantee the quality of 'time sensitive' services compared to using the 'best effort approach' of treating all data traffic in an equal manner.

Extra confusion in this discussion is added by the - not always properly used - distinction between services and content, especially when time critical services, for instance two-way communication like video conferencing, are discussed. Confusing terms such as 'traffic time-sensitive content'[2] are used in situations were a limited delay does actually not have disadvantages for the user. Too often people discard the fact that audio and video data are in many cases delivered in a faster speed than they are played back [3]. Together with the smart use of buffering this solves many of the jitter and latency issues in one-way traffic. Therefore the 'time sensitive' argument is only valid for services where even a small delay is

unacceptable and buffering is not an option, but it is less valid for services like video distribution where buffering is an adequate solution.

It is therefore important to distinguish the relevant services or service categories, their providers and their subsequent relations, in order to deduct which parties are involved with the neutrality of the network and the transparency of that related to users and service providers. For this a schematic overview of the parties, and their relationships, is needed.

3.3 Costs and benefits

Using network traffic shaping equipment it is possible to achieve QoS prioritising via standardised Type of Service bits, but this is only effective if this is done end-to-end, in all segments of the network between source and destination, or at least in all segments that are insufficiently dimensioned. Also here, the chain is as weak as its weakest link.

Discrimination based on source or destination address is relatively easy, but deep packet inspection for protocols or content to be able to distinguish different types of traffic is relatively costly in terms of processing power and affects the roundtrip delay. Besides, when the intention of a technical solution is to prioritise time-sensitive traffic, the effective deployment of the solution will be difficult, as distinguishing traffic that is really time sensitive from traffic that only appears to be time sensitive but is in reality 'best effort' is not always simple. Port numbers and protocols are often reused as carriers for traffic that they are not originally intended for. Prioritising one type of traffic by definition means giving other types of traffic lower priority, thus creating an incentive to disguise traffic as 'high priority'.

In the case of shared (access) media, variations in available bandwidth for a user can only be limited by actively limiting used bandwidth of (other) users, or by sufficiently over-dimensioning the network, the costs of which varies per network type and legacy situation. The possibilities to upgrade a specific type of network are sometimes limited by technical capabilities, or by boundaries set by regulators and standards. In that case it may be necessary to either use another type of access network able to provide the higher capacity needed, or to extend the backhaul network and split up the existing access network in segments, providing each segment with a separate high bandwidth connection to the backhaul network. , Both of these solutions are costly.

Over-dimensioning and multiple peering agreements that allow alternative traffic routes are mostly sufficient to cope with variations in bandwidth use in the core network. But here is also an economic issue due to peering SLAs etc. that are enforced by border gateway control between the domains of different network owners.

Incentives for stakeholders to promote or preserve net neutrality are the general global accessibility of services and content. This is especially important for globally operating companies (and the Internet has become the means by which also many small companies are now globally operating).

Incentives for stakeholders to break or decrease net neutrality are the technical and economical advantages of personalization, increasing context awareness and other forms of differentiation of offered services and content. Prioritising certain types of traffic over other types of traffic may increase efficiency of networks (less over-dimensioning needed), and enables charging for specific service or content types (as part of or additional to the offered Internet access). It can also be a means to protect revenues from already existing services such as telephony and cable television.

4 Towards workable models and definitions

To get a grip on the complexity surrounding the net neutrality discussions two steps have to be taken:

- We need to avoid overly simplifying the problem area, using a model that can be used to map the service playing field and the parties that operate in it in a framework that encompasses several aspects of net neutrality. The scope of the model is deliberately wider than only ‘the net’, as the end user perceived neutrality of the net is also defined by the neutrality of the (most important) services that are offered over the net.

- We need a workable definition of the term ‘neutrality’, since this concept can be associated with access and accessibility, but also with general and vague associations with speed or relative quality of Internet services.

4.1 A practical model: visualising spheres of influence

First of all we can distinguish different orthogonal aspects that are associated with the term net neutrality. Besides the aforementioned roles of the end user, it is important to get insight in the sphere of influence on different functional levels and on different geographical levels. For every combination of these two aspects, if there is little (opportunity for) competition, there is a possibility that a *control point* emerges: one or a few stakeholders dominate a certain aspect that eventually has end-to-end implications, and can influence or limit net neutrality. An example of such a distribution can be found in figure 4.

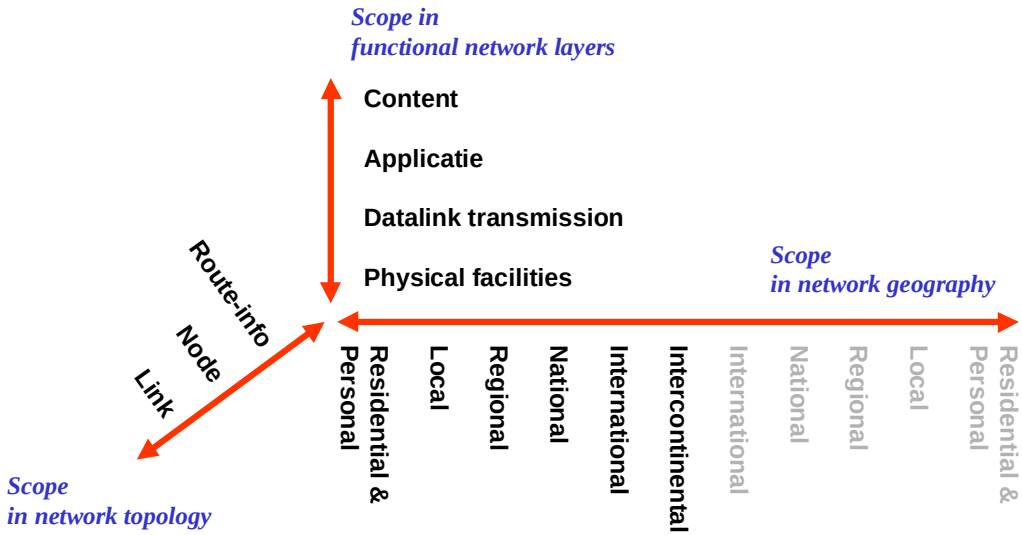

Figure 4: Axes of Influence for communication- and distribution services

Control points emerge when a party has exclusive control over an essential network function. Network geography can be divided in areas each with specific functional scope:

Residential & Personal: For functions having personal or residential scope, such a provided by residential equipment and software or end-user equipment, software and content.

Local: For functions having a local scope, such as related to the access network that directly connects the residence of the user or the mobile user (last/first mile, local loop) or locally generated content.

Regional: For functions having a regional scope, such as provided by the regional network, connecting the local access network to the core network of an access provider, and regionally operating applications

National: For functions having a national scope, such as provided by many Internet access providers, but also by web content and services aimed at national markets

International: For functions having an international scope, such as provided by internationally operating Internet content and application providers.

Intercontinental: For functions having a world-wide scope, such as provided by transatlantic backbone providers.

This division can be further simplified or detailed for a certain case.

For every communication or distribution service those functions can be categorised in:

Links: physical or logical connections between sources or destinations of data (example: a DSL connection)

Nodes: physical or logical transhipment of data (example: a router)

Routing Info: Information necessary to get data from one place to another (example: a number directory, but also DNS the information that maps site names to IP addresses)

It is very important to focus on all three functions and not limit the scope to one or two. For instance mechanisms such as DNS are essential and more vulnerable then people realise [13].

Using the axes of influence and this categorisation, for a specific service (like web browsing or IP television) an overview can be made of functions necessary for the deployment of the service. A (simplified) example of such a model for the service 'web browsing' can be found in Table 1.

Table 1: Control points matrix: functional break up of market parties, functional layers and position in network geography (Informative, the functions in the cells are not limited to the mentioned examples)

	Residential & Personal	**Local**	**Regional**	**National**	**International**
Content	web page		DNS entries	news portal	Distributed database
Application	web browsers, DNS caching		DNS lookup	Internet hosting	distributed search engine

Network	IP subnets	edge & core routers	ISP	ISPs	transit providers
Datalink transmission	DSL modem, cable modem	cable			transatlantic backbone providers
Physical Facilities	cables, ducts	cable	dark fiber	dark fiber	sea cables, satellite connections

In the resulting overview the stakeholders that offer these functions, and also the money that flows between the companies, can be depicted, thus showing the spheres of influence of the stakeholders, and the existence of power concentrations. In the case of an Internet service in competition with another ('old') service the model can be made for both cases ('old' and 'new'), giving insight in where, why and how market parties may try to influence Internet traffic.

In cases where it is unclear what could be the costs and benefits, a comparison can be made of the costs and benefits in different scenarios. For instance a comparison can be made of one scenario in which a company tries to overcome capacity problems by overprovisioning, and another scenario in which the same company uses prioritising traffic based on type or source.

4.2 Example cases

A model such as described above can prove a good tool to map and investigate current and future net neutrality threats and the effects of different forms of market transparency. Two short examples are given to show how the tool can be used to discuss different aspects of the problems mentioned above. The first example describes a streaming video service, the second example describes mobile apps.

An example is the case of a user watching a CNN news video (see Table 2). In blue the spheres of influence of companies are depicted, and with red arrows the money flows between the companies. We see that the end user pays for Internet access to the access provider. This company pays the DSL provider, which sometimes pays other parties for providing the physical infrastructure. Indirectly the end user pays CNN (by watching ads, for which the advertiser pays). CNN pays its access provider. Finally the user's ISP pays the Transit Provider for the use of the backbone. Microsoft is mentioned in this overview because this company provides the used Web browser and operating system of the end user. Web application builders have to be aware of the sensitivities of Microsoft end user products to be able to provide a proper working service [14].

Table 2: Example of a Control Point matrix for a streaming video service with spheres of influence: breakdown of market parties, functional layers and position in network geography, with included parties and money flows from the example in the text

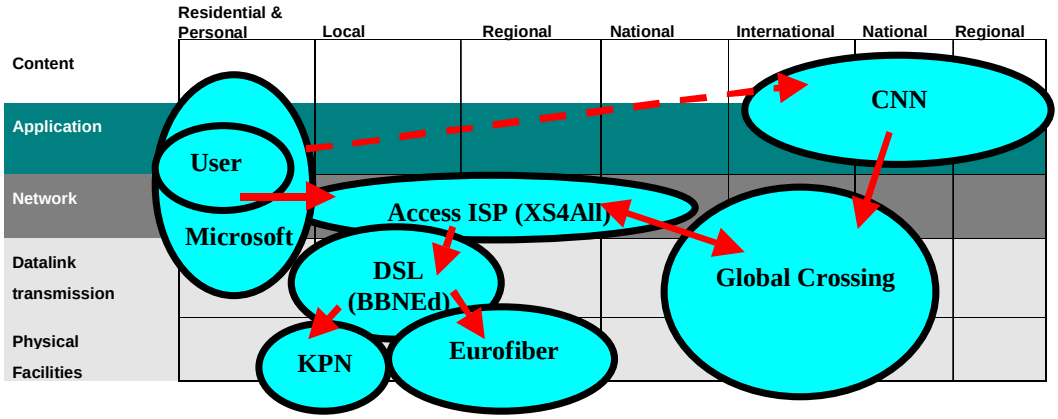

This is only one specific situation, that we can compare with parties and money flows that are related to the receiving of the same CNN item via cable television. By constructing dedicated examples we can quickly make comparisons, and gain insight in technical and economical motives of the different stakeholders.

Another example is depicted in Table 3. In the mobile telephone industry other types of Operating Systems are used than generally used for laptops and PCs, with other characteristics. Mobile operators are trying to protect their revenues from selling voice connections and sending and receiving SMS messages. This makes it especially tempting to not encourage using similar services over the Internet connection offered on mobile phones. Also the operators are involved in business regarding selling of ringtones and special apps, types of services that are not always easy to install by the end user when downloaded 'directly' from the Internet.

Table 3: Example of a Control Point matrix for mobile apps with spheres of influence: break up of market parties, functional layers and position in network geography, with included parties and money flows from the example in the text

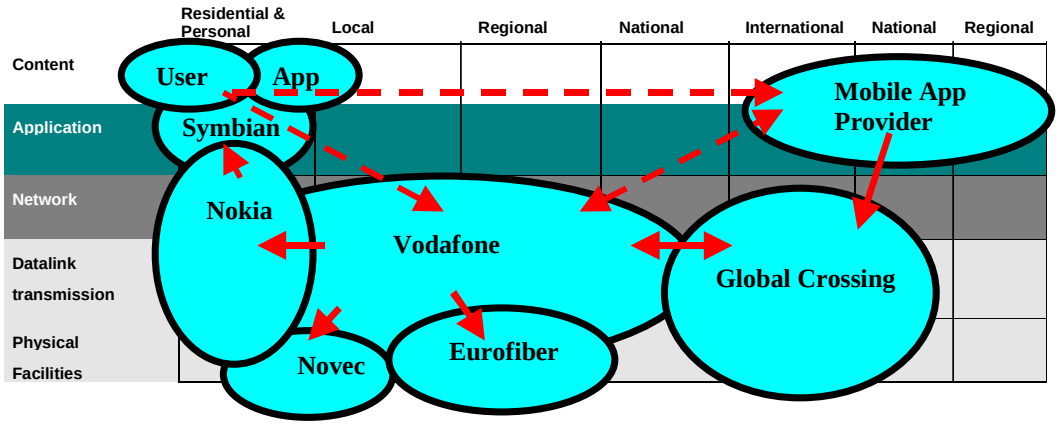

Also here other diagrams can be Similar examples can be made for other stakeholders and dependencies, as a starting point to get insight in technical and economical motives of the different stakeholders.

4.3 A workable net neutrality definition

As shown in chapter 2 and 3, the meaning of the term net neutrality varies. It depends on the viewpoint and role of the stakeholder, on which functional layers taken into account, on the position in the network geography taken into account, and even on the functionality in the network topology that is discussed. Additionally, definitions vary based on the opinion on neutrality with regard to differentiation of quality of service levels.

From a regulatory perspective the definition [12] given by the Trans Atlantic Consumer Dialogue (TACD) is interesting because of the consumer perspective, both as content requesters as well as content providers. It covers the currently most important aspects in the Control Point Matrix. For obvious reasons it focuses on the consumer role of the end user, but with very little modification it can be generalised for the end user in different roles. The enhanced definition is given below:

In a neutral Internet, end users:

- have the right to attach devices of their choice; (user domain, physical layer)
- have the right to access content, services, and applications of their choice; (all public domains, datalink, network and application layer)
- have the right to provide content, services, and applications of their choice; (all public domains, datalink, network and application layer)
- have the right for their access to be free from discrimination according to source, destination, content, or type of application.(all layers, all domains)

Simultaneously, in a neutral Internet, ISPs and public communications networks:

- do not unfairly block content, applications or devices;
- do not deliberately degrade access for content, applications, or devices;
- do not prioritise data according to its source or destination;
- do not discriminate against particular providers of content, applications, services, or devices.

Additionally the above mentioned prohibitions also apply to end user hardware and end user software that claims to provide Internet access (such as web page browsers), unless this is specifically selected as a special service by the end user.

If we follow the Control Point Matrix for every topology function (link, node and routing info) we see possible other threats, for instance the potential monopoly of one party in the world wide search engine market.

5 Conclusions and Recommendations

Net neutrality is a complicated and multidimensional subject, and many aspects are not sufficiently recognised in many of the discussions surrounding this subject. In this article an attempt is made to untangle the related issues by distinguishing different functional layers, topology functions and network geographical scopes. It is very easy to step into the pitfall of solving a problem on a level where there it does not exist.

Three main threats of net neutrality can be distinguished:

1. Market dominance on a control point, at a certain geographical domain on a certain functional layer.
2. Insufficient transparency regarding offered connectivity, services and service capabilities
3. Insufficient transparency regarding the economical relations between stakeholders providing connectivity, stakeholders providing services and stakeholders providing certain content.

Net neutrality threats are not limited to the home and access domain. The model presented helps to determine the most effective combination of functional and geographical layers on which to implement regulation and standardisation. Since most Internet and telecom services span more than one country, regulation on a multinational level is to be preferred.

When proposing regulatory solutions for one of the problems that are grouped under the term ‘net neutrality’, it is important to clearly state the specific goal of the solution. Only pointing

vaguely to the term net neutrality as a goal is insufficient. It may have advantages to investigate whether it is possible and necessary to define a basic level of Internet access for end users (and come up with a general definition of what Internet access actually is), and what transparency is necessary for service levels on top of that.

A potential danger of regulation is that it may lead to unnecessary complicated and costly solutions for situations that are not problematic. To a certain extend the nature of the Internet will lead to minor fluctuations in quality, depending on overall and local usage, position in the network topology etc. Ex post regulation may be better suited to safeguard the spirit of net neutrality than ex ante regulation.

The Internet has facilitated and created the preconditions for the conception and evolution of innovative new services by providing a world-wide open basic communication platform. Further investigation is necessary how to ensure that this aspect of the Internet is not lost in perfecting it. Monitoring and - when necessary - regulating with regard to dominance of a limited set of stakeholders on control points should be strongly encouraged to prevent possible misuse of a dominant position.

6 Acknowledgements

The author would like to thank Paul Brand, Alexander ter Haar and Hendrik Rood of Stratix Consulting for their valuable input and discussions.

7 References

[1] TILEC, Universiteit van Tilburg, Network neutrality and the nature of competition between network operators, 11 July 2007.
[2] GÜNTER KNIEPS and PATRICK ZENHÄUSERN, The Fallacies of Network Neutrality Regulation, CNRI, 9(2), June 2008
[3] ANDREW ODLYZKO, The delusions of net neutrality, School of Mathematics, University of Minnesota, August 2008
[4] CHRISTIAN SANDVIG, Network Neutrality is the New Common Carriage, University of Illinois at Urbana-Champaign
[5] "Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway", ANDREW TANENBAUM.
[6] TRANS ATLANTIC CONSUMER DIALOGUE, Resolution on Net Neutrality, March 2008
[7] Citizens, Communications and Convergence, OFCOM DISCUSSION PAPER, July 11th, 2008
[8] MARCEL VAN KUIJK, Recht op toegang tot Internet, University of Tilburg, the Netherlands.

[9] ANTON EKKER, Anonimiteit en uitingsvrijheid op het Internet; het onthullen van identificerende gegevens door Internetproviders, Mediaforum 2002-11/12, p. 348-351
[10] CHRIS ANDERSON: The Long Tail
[11] RECOMMENDATION CM/REC(2008)6 of the Committee of Ministers to member states on measures to promote the respect for freedom of expression and information with regard to Internet filters
[12] TACD Trans Atlantic Consumer Dialogue, "Resolution on Net Neutrality", DOC No. INFOSOC 3608, March 2008
[13] BLOG WIRED MAGAZINE: "ISPs' Error Page Ads Let Hackers Hijack Entire Web, Researcher Discloses", Brian Singel, http://blog.wired.com/27bstroke6/2008/04/isps-error-page.html
[14] Wikipedia on standards support of Microsoft Internet Explorer 7, http://en.wikipedia.org/wiki/Internet_Explorer_7#Standards_support
[15] My Internet is not the same as Your Internet, Mashable, http://mashable.com/2009/04/27/my-internet-your-internet/
[16] http://www.ispam.nl/archives/9861/upc-erkent-afknijpen-usenet-verkeer
[17] EUROPEAN COMMISSION, Questionnaire for the public consultation on the open internet and net neutrality in Europe, http://ec.europa.eu/information_society/policy/ecomm/doc/library/public_consult/net_neutrality/nn_questionnaire.pdf