

Traffic Management and 'net neutrality'

Response of Everything Everywhere to the European Commission Consultation

About Us

Everything Everywhere Ltd is a new joint venture, 50% held by each of Deutsche Telecom and France Telecom, and which operates the T-Mobile and Orange brands in the UK. Everything Everywhere serves approximately 30m mobile customers and is one of the main UK fixed broadband providers also.

Executive Summary

The UK market for Internet access is highly competitive. At a retail level the number of ISPs, connection speeds and broadband penetration have increased rapidly, in both fixed and mobile, while at a content/service level, more and more public and private companies serve consumers through the Internet, with high levels of innovation and take up. This is reflected by ever increasing numbers of broadband subscribers using the Internet for an increasing array of services.

With high levels of consumer take-up and competition among providers, the UK market for ISPs is focussed primarily on the price and speed of access. This requires that network operators both invest heavily in network capacity, and optimise network efficiency, particularly at peak times. This is because ISPs must serve the conflicting demands of high speeds and low costs: consumers demand high speeds (requiring network investment in capacity), at low cost (requiring the limitation of capital expenditure). With both speed and costs driven by the level of peak time capacity provided, networks must optimise the balance between peak time speeds and the impact on retail charges to arrive at a competitive position. The consequence is that, in the absence of an incentive on the companies generating Internet traffic to optimise this (i.e. some form of price signal), there must be some measure of network management in order to prevent peak time speeds and reliability suffering as a consequence of network congestion. Typically this entails the prioritisation of the traffic most valued by the majority of customers at peak times. The important point is that, in the UK, network management measures are driven directly by the need to meet consumer demand for low cost, high speed access.

To date the Internet has flourished precisely because of the lack of regulation, and absent evidence of market failure or demonstrable consumer detriment it is unclear on what basis regulation should be imposed given the risk of constraining further innovation. There is a considerable danger that regulation to enact particular rights or freedoms in fact serves only to freeze the form of the Internet at a particular point in time. In the longer term this will act only to restrain new developments and to serve specific interests which benefit from the current approach to content distribution. Consequently there is a danger that new "rights" in fact become restrictions, limiting both consumers and suppliers by forcing particular forms of service / payment / engagement / relationship and preventing new forms of each.

Hence, it seems appropriate to allow users and providers to determine how the Internet develops. This does not mean that no regulation whatsoever is required, but that this need not necessarily be Internet specific. Rather, the question is to ensure transparency of the market to ensure that it remains competitive, innovative etc, and that market mechanisms are enabled to allow informed customers to determine what services/innovations flourish and fail. Accordingly, net freedoms are best served by ensuring that the market works by ensuring that consumers are able to make well informed choices and compare competing offerings.

Finally, regulation of the Internet cannot be driven through the ISP layer alone. Not only is this market competitive, but it is one of only a number of layers in the delivery of internet services: many aspects of performance and availability are determined by the content/application providers etc. So long as ISPs are not uniformly blocking particular content/services across

their range of tariffs etc then there is neither any need nor basis for intervention, which could significantly reduce efficiency, innovation and therefore consumer benefit.

Question 1: Is there currently a problem of net neutrality and the openness of the internet in Europe? If so, illustrate with concrete examples. Where are the bottlenecks, if any? Is the problem such that it cannot be solved by the existing degree of competition in fixed and mobile access markets?

Overview of the UK market

In summary, Everything Everywhere believes, supported by Ofcom's findings in the UK, that the UK broadband markets are competitive. We believe that the UK market is delivering high quality, low cost and widely available services and that it is serving the consumer interest. Although networks face congestion from time to time, this does not reflect a problem of net neutrality, nor lead to one. To the extent that any network were to begin to constrain neutrality and openness in a manner not acceptable to consumers they would rapidly lose market share. Hence, the competitiveness of the UK market, while not guaranteeing perfect outcomes in every scenario, ensures that there is no particular need for additional regulation.

In the UK there are a large number of competing ISPs and networks. Competition in the market to provide fixed internet services has rapidly increased over the last 10 years, with a number of ISPs providing internet access as part of a package of home utilities (satellite, cable and terrestrial TV, fixed and mobile telephony etc). Among mobile operators, all the MNOs and a large and increasing number of MVNOs offer data tariffs with a high level of competition on access and equipment charges, speed, availability and reliability. Indeed, many ISPs (both fixed and mobile) have provided internet access below cost, to leverage existing infrastructure or to create multi-play offerings in order to increase their competitiveness and customer loyalty. In parallel both fixed and mobile networks have invested heavily in providing increasing speeds of access at lower cost.

As a consequence the UK market for broadband access is characterised by the large number of competitors and a focus on the price and speed of access:

Figure 1: Price of mobile broadband

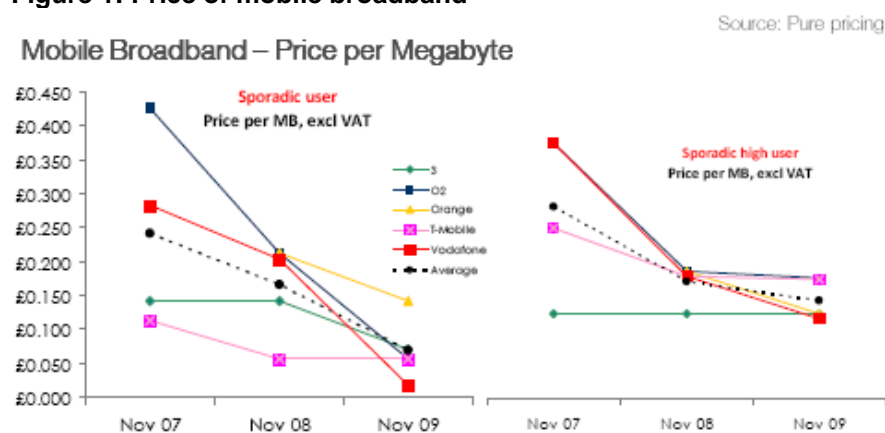



Figure 2: Headline broadband speeds

Figure 5.9 Non-corporate broadband connections, by headline speed

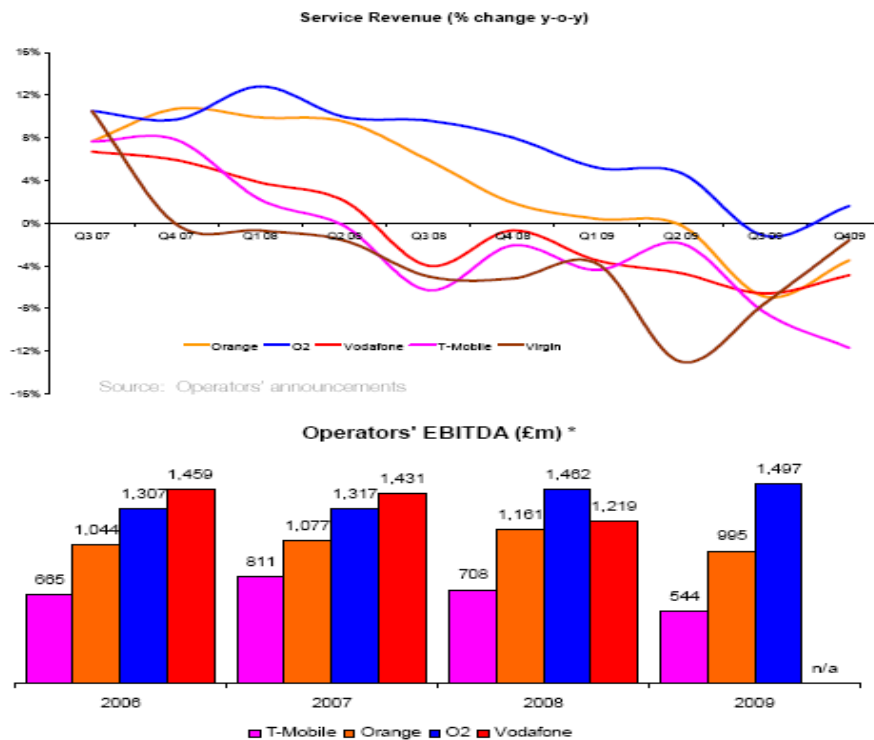


Source: Ofcom / operators
 Note: Includes estimates where Ofcom does not receive data from operators; excludes mobile broadband connections

As can be expected, the consequence of this high intensity of competition is a high level of product and service innovation, decreasing costs, increasing availability and robust customer satisfaction (according to the Ofcom Communications Market Report 2010, 83 and 90% of customers are either satisfied or very satisfied with mobile and fixed broadband services respectively).

Among operators this is further corroborated by high levels of network investment and low profitability - the competition for market share has exercised considerable downward pressure on prices, as illustrated by figure 3 below.

Figure 3: profitability among mobile operators



Unlike the other top operators Vodafone does not report EBITDA to the end of calendar year but from April to March

It is notable that the profitability of UK mobile ISPs is the lowest in Europe.

The supply chain

While it is the ISP that sells Internet access to the consumer, the ISP itself is not wholly in control of the service it offers however. Hence, even if there were an issue in respect of net neutrality and openness, it does not follow that this arises as a consequence of the action/inaction of ISPs.

As is common knowledge, the Internet is a network of networks on which content and services are stored and served. The reliability or speed (for example) at which a particular item of content is delivered is therefore contingent on, but not entirely determined by the ISP: delivery of any content or service is determined by the "weakest link".

The nature of the "weakest link" may be various: it could be the CPU speed of the server on which the content is held, the bandwidth of its connection to the Internet, whether the content owner has contracted for a particular quality of service etc. Hence, performance can be varied in various ways at various points in the supply chain – by various players.

Most significantly performance and availability will be heavily influenced by the level of investment made by:

- the content/service owner in delivery and distribution: i.e.
 - the extent to which they have distributed their content/service to local servers, or whether it is hosted only in their local ISP
 - whether they/their ISP rely on peering for delivery (i.e. best efforts) or are paying for transit (and therefore a particular QoS)¹
- the underlying network
 - ISPs are contingent on the underlying network installed by their wholesale supplier. For example, in the UK many ISPs rely on the network supplied by BT Wholesale. The level of investment and service installed and supplied by the underlying national fixed telephony network(s) will determine what service or services an ISP can deliver. Similarly, the cost of capacity on the wholesale network will determine the speed and structure of retail tariffs, and in turn whether traffic management is required in order to avoid raising costs.

It follows that any examination of net neutrality must examine the wider web ecology, since any narrower examination will inevitably:

- lead to incomplete and partial conclusions, and proposals by focussing only on one level of the supply chain
 - As noted above, not only do content/applications providers etc significantly influence the nature and performance of the Internet, but they have a significant place in determining the efficient use of the networks and efficient investments in networks. As the volume of traffic generated increases the need for efficiency increases. Investing in capacity alone will not “solve” congestion (even with network management solutions) since any new capacity is automatically consumed by additional traffic. Consequently, the current business model of peering and free but growing imbalanced traffic will have to evolve.
- distort the market
 - this is particularly significant in the case where many ISPs provide parallel services which may compete with those provided by other companies over the Internet. For example, an ISP subject to strict restrictions on the use of data may not be able to compete equally with an over the top operator (“OTT”) subject to a wholly different regime.

Similarly, the principles of net neutrality can be equally influenced not only by ISPs but by the content, application providers, international carriers and underlying network providers themselves, which are equally well placed to drive efficiency and otherwise control the distribution of traffic/content/services etc. For example, many previously free services provided on the Internet are now subject to subscription (e.g. the newspaper The Times), or may not be available outside particular jurisdictions (e.g. the BBC restricts the distribution of its free online content to particular territories and resells this in others).

However, the point here is not necessarily that these other actors should be regulated, but that regulation is inherently dangerous in a market which is as complex and rapidly developing as the Internet. Absent any evidence of consumer harm the market should be enabled to develop those services, structures and assets as will best serve the consumer.

The consumer offer

It follows from the above that the market for Internet access and services/application etc, is competitive, with a high degree of service innovation etc (e.g. the BBC iPlayer and the plans for YouView - a cross platform IPTV service).

¹ The clearest single example of the above is inevitably Google, which:

- has invested heavily in localised hosting of content/services
- makes additional investment in QoS for particular services over others
 - as noted above, Google pays transit in some countries in order to guarantee the QoS for its most profitable services, while relying on peering for others
- intends to spend up to \$1.6bn on the installation of fibre networks in trial cities providing speeds of up to 1GBps (http://www.businessweek.com/technology/content/feb2010/tc20100211_381119.htm)

While it must be acknowledged that traffic management takes place, this is not a “bad” thing. Rather, traffic management enables networks to provide a high speed low cost service efficiently to the consumer, and in effect reflects the consumer demand for these qualities. Traffic management is therefore a pro-consumer exercise.

It follows that not all services/applications etc are available on the same basis at all times, but that these are prioritised at times of congestion in order to maximise the experience of the majority. While there are some ISPs that will, on some tariffs or services (e.g. mobile) make some services available that are not available on others (e.g. VoIP not available on entry level tariffs), this blocking is typically not universal, either across that operator’s own offerings or that of the industry. In this form, such traffic management enables the creation of differentiated offers without precluding a particular OTT services (for example) from accessing, or being accessed by, consumers. In a competitive market there is no basis to determine on a top down basis what consumers should or must have access to: they can determine this through exercising choice.

In summary, neither the capacity of the Internet, nor the capacity of consumers to pay for capacity etc, is infinite. In these circumstances, traffic management is a necessary tool to enable the delivery of the best service to the most consumers. Moreover, it is not ISPs alone that determine the service delivered to the consumer, which can be influenced at a number of levels in the supply chain, including by the content/service providers themselves (see further below). None of this however is a de facto negative characteristic, but, so long as the market remains competitive, is a function of delivering the services requested at the most appropriate/lowest cost.

Accordingly, there is no net neutrality “problem” in the UK, nor likely in Europe (where national markets will face the same constraints and access the same global network of networks). While not all services etc may be delivered in the same manner at peak times, this is a reflection of the constraints on capacity and demands of consumers. Furthermore, the innovation and efficiency that traffic management allows is enabling the development of further services at lower cost.

In summary, so long as the market remains competitive there is no reason to compel uniformity on the services offered to consumers.

Question 2: How might problems arise in future? Could these emerge in other parts of the internet value chain? What would the causes be?

As stated in the previous answer, there is at present nothing to suggest that the market is not functioning effectively and competitively in response to consumer needs and requirements, nor any indication that it will not continue to deliver. (This is also evidenced by the lack of complaints brought to the attention of Ofcom).

Everything Everywhere furthermore believe that consumers will benefit from enhanced transparency measures and is actively engaged with Ofcom to ensure that our customers’ needs will be met through whatever provisions are put forward (please see question 5 for further detail).

There are scenarios where problems could emerge, primarily linked to those in which the underlying content providers/aggregators etc develop and abuse a dominant position. A significant proportion of traffic and services are generated via a relatively small number of players. Similarly, the provision of applications via the internet is, in some respects, subject to the policy decisions on availability etc of a limited number of core platforms/portals. While it remains unclear how the market will develop over time, it does not appear that there will be any particular novel problem or cause to it, but that any that do develop will derive from a lack of competition.

Question 3: Is the regulatory framework capable of dealing with the issues identified, including in relation to monitoring/assessment and subsequent enforcement?

The regulatory framework has three distinct strands that relate to this issue, around 'open Internet', 'transparency' and 'Quality of Services'. It is worth noting that the framework is still in its transposition phase in the UK and there remains a lack of clarity on what its final implementation will look like and mean for the UK. Our initial view however is that we believe that there is no requirement at this stage to further extend regulatory intervention or provide additional powers or tools in response to an issue where there is no immediate or identifiable problem. Furthermore, until the framework has been implemented in Member States, and been given a chance to effectively function, it would be premature to impose further regulatory burden or extend legal uncertainty for operators.

Transparency

We believe that of the three areas identified in the regulatory framework, the most appropriate strand for empowering and benefiting consumers lies in providing them with important and relevant information that will prove to be both meaningful and relevant in their decision-making. Everything Everywhere is actively engaged in the UK discussion currently being carried out by Ofcom and we seek to ensure that any transparency impositions are not overly prescriptive and consequently lead to a reduction in competition (by harmonising the points of competition) or innovation (by preventing the presentation of tariffs that do not conform to the parameters of an agreed presentation).

While an adequate remedy, Everything Everywhere however does not believe that there is a market failure as regards transparency etc that requires remedy. As set out above, Everything Everywhere considers that there are significant potential dangers in any mandated form of transparency and that this is best left to the market until such time as there is a clear case for regulatory intervention. More importantly however, there are significant incentives on competitors to increase transparency in order to create differentiated tariffs, enable more evolved traffic management, and differentiate themselves from each other. Hence, the combined factors of a competitive market and the need for traffic management will themselves drive transparency.

Quality of service

Absent a finding that the market is not working in the interests of consumers or is otherwise not competitive, Everything Everywhere does not believe that there is a basis for imposing QoS obligations. Indeed, as set out further above, particular caution should be exercised in regulating a fast moving market in which there is high levels of innovation, investment and consumer satisfaction. It is therefore appropriate to either not regulate in this manner, or to try a number of less intrusive alternatives first.

Everything Everywhere does not believe that there is a basis for additional regulation as there is no current market failure. Regardless however, transparency on QoS etc is a significant and increasing feature of the market since it enables

- consumers to choose between (and therefore demonstrate demand) different services; and
- competitors to differentiate themselves and their offers – fostering innovation and further investment.

Everything Everywhere is therefore keen to ensure that the market continues to function in the interest of consumers and competitors and to reinforce transparency consistent with being able to offer differentiated services.

Everything Everywhere believes that this will be sufficient to address any potential concerns. In particular, transparency will enable consumers to demonstrate what services or tariffs there is demand for: for example, if consumers demand mobile VoIP then transparency that this is available on [x] tariff and not [y] will enable competitors to provide, and consumers to choose, tariffs including VoIP². Similarly, if consumers decide that they are willing to pay more for

² As set out further above, there may be other parameters to such competition, such as guaranteed QoS, or "inclusive minutes" (e.g. VoIP not counting towards download limits etc).

completely unlimited/unmanaged Internet access then they will opt for those tariffs that make clear that P2P traffic is not throttled etc.

Furthermore, imposing particular QoS creates its own danger of consumer harm, since it will (as set out above) harmonise points of competition and reduce the scope for innovation. Everything Everywhere believes that minimum QoS therefore represent a potentially significant consumer harm in themselves, both generally (by limiting traffic management / requiring higher retail prices to permit additional capacity investment) and in respect of particular services (e.g. smart meters, which may specifically only require low bandwidth connections).

That being said, national regulators are empowered under the regulatory framework to impose minimum QoS on networks. We believe that such an imposition should be considered only if it is established that there is considerable consumer harm being caused by low minimum QoS. Given the current (and rapid) inflation of QoS, this would likely be analogous to a finding that the market is not competitive. Even in these circumstances, given the need to ensure a proportionate approach it should also be demonstrable that there is no less distorting or interventionist approach that could address the issue in the circumstances.

Furthermore, the regulatory framework refers only to the 'undertakings providing public communications networks' as regards setting minimum quality of service requirements. This approach would ignore the role played by other actors in the internet value chain who are equally able to directly influence the QoS delivered to the end-user (as stated previously in the response to question 2).

Open Internet

As will have been recognised by the Commission, there are several interpretations as to what is meant by the term 'open Internet'.

There are no incentives for unfair discrimination that are particular to the Internet. Competitors will continue to compete vigorously, with those that fail attributing a measure of blame to the market, its regulation and their competitors.

So long as the market is competitive, and there is transparency to consumers about what services are enabled/prioritised etc, discrimination between companies is an inherent part of the competitive process. The right to protect investments and to innovate around different tariff structures etc ensures continued investment and innovation, enabling competitors to continue to develop new products and services. A measure of discrimination is therefore inherently pro-competitive and pro-consumer. For example, discrimination underlies exclusivity, which is the basis for many new (pro-competitive) product launches.

As stated above, so long as customers are informed of their options, a range of tariffs available is a clear demonstration of a competitive market and consumer behaviours and demands will drive consumer offerings.

Question 4: To what extent is traffic management necessary from an operators' point of view? How is it carried out in practice? What technologies are used to carry out such traffic management?

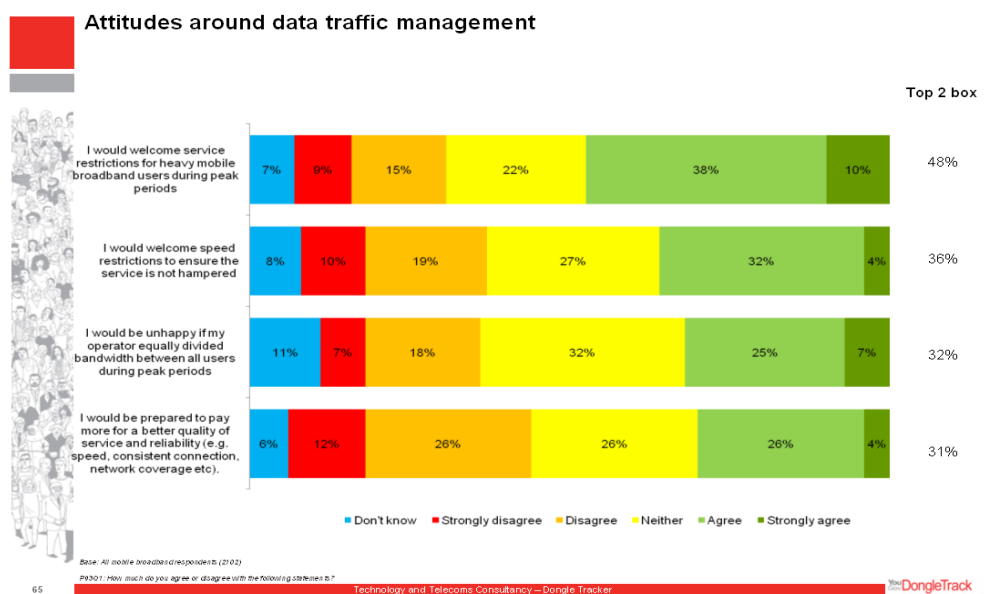
In response to this question, several points are worth bearing in mind:

- different types of internet use require different QoS (i.e. video requires higher QoS than email);
- network capacity has a cost, and networks must be dimensioned to deal with peak demand;
- networks (mobile in particular) are increasingly facing congestion at particular bottlenecks;
 - this is particularly the case in mobile where an increasing number of urban cell sites regularly reach congestion at the air interface. The high cost of high capacity backhaul is a growing problem also (e.g. the BT GigE product).

- traffic management is therefore required in order to ensure that a small minority of heavy users / non-time dependant services, do not prevent/unnecessary degrade performance for the majority;
- traffic management therefore enables more efficient network utilisation through;
 - reducing the total network capacity required where traffic management can reduce peak demand – enabling networks to offer lower prices by reducing their network capacity requirements;
 - maximising the capacity available for the majority; and
 - enabling the provision of a wide range of services to all customers, which may not be possible where all services had to be provided with limited resources: i.e. HD video may not work / be commercially viable at peak times where there is no traffic management.

Much of the above is summarised by the findings in the latest YouGov survey on mobile broadband:

Figure 7: consumer attitudes to traffic management



i.e. consumers would much rather accept traffic management than higher prices.

There will be specific instances where certain services/applications may be impacted negatively (albeit potentially only at certain times of day/within certain parts of the network) by traffic management. For example, there may be applications for P2P file sharing that do not evolve/evolve less quickly because P2P traffic is more frequently subject to traffic management at peak times (owing to the disproportionate use of capacity for a relatively small number of users).

However, overall, traffic management drives innovation, for two reasons:

First, as set out above, there is a capacity driven need for some network management in order to ensure that networks can physically deliver the service demanded by the majority. In congested areas consumers would be prevented from using HD video (for example) where heavy P2P traffic is unrestrained. There is far greater consumer demand for and innovation in new Internet delivered video content services (e.g. the BBC iPlayer) than in the delivery of existing content directly between users (e.g. P2P).

Second, traffic management may be necessary to ensure that networks can deliver a business case to evolve new services and tariffs. For example, as set out above absent guaranteed QoS it may not be possible for HD online gaming to be delivered (since a “best efforts” service will not drive sufficient customer satisfaction to enable a business case reliant on continued subscriber revenues).

Finally, it should also be noted that some specific services require specific traffic management such as e-health or e-mobility services or machine to machine communications.

Only if network capacity was infinite would there not need to be any network management. However, not only would the provision of such levels of capacity be inherently inefficient, but it would have a direct impact on price, which would in turn inevitably impact negatively on customers. In effect, the ideal of net neutrality is inherently conflicted: while the principle of a “free Internet” is ideal, it would in fact result in more expensive Internet. Alternatively, absent a price rise, it would result in network congestion and a “net freedom” to enjoy a congested service. Furthermore, the provision of additional capacity will not curtail congestion alone: Internet protocols are designed to automatically increase traffic generation until congestion is reached. Internet traffic is like a perfect gas: it occupies all the capacity which is provided. Increasing capacity alone cannot prevent congestion.

Moreover, traffic management may be inevitable where capacity is unavoidably constrained. This is particularly the case for mobile networks where spectrum scarcity, cell tower planning constraints (and installation cost) and the availability of economic high capacity backhaul mean that cell congestion cannot easily be avoided through capacity expansion. Although there are various potential remedies to increase total available capacity, these are neither short term in implementation, nor long term in solution: capacity expansion is expensive and time consuming and (given demand growth) can only postpone congestion rather than avoid it, since traffic will always expand to consume available capacity (see above). Hence, active management of demand for capacity is required. As noted above, in time this may require some form of economic incentive for efficiency at a “wholesale” level also if the long-term sustainability of the open internet is to be ensured.

For ISPs network management is therefore key to the provision of a competitive service. With costs significantly determined by the cost of peak capacity, managing peak demand has become a critical element of competitive success. Other than by raising prices, it is simply not economic to provide all services on a best efforts only basis and for them all to function adequately: at peak times network congestion would mean that services requiring low latency etc would simply not function while other uses, such as general web browsing, may become unacceptably slow.

Everything Everywhere therefore performs traffic management to ensure load management at peak times by filtering traffic by protocol. This is only implemented when the network reaches congestion however.

To the extent that there is a concern regarding traffic management by networks, this should be seen as a positive rather than as a consumer harm: traffic management enables higher quality of service to the majority at lower cost.

Question 5: To what extent will net neutrality concerns be allayed by the provision of transparent information to end users, which distinguishes between managed services on the one hand and services offering access to the public internet on a 'best efforts' basis, on the other?

Transparency on QoS is a significant and increasing feature of the market since it enables:

- consumers to choose between (and therefore demonstrate demand) different services; and
- competitors to differentiate themselves and their offers – fostering innovation and further investment.

Everything Everywhere is therefore keen to ensure that the market continues to function in the interest of consumers and competitors and to reinforce transparency consistent with being able to offer differentiated services.

Everything Everywhere believes that this will be sufficient to address any potential concerns. In particular, transparency will enable consumers to demonstrate what services or tariffs there is demand for: for example, if consumers demand mobile VoIP then transparency that this is

available on [x] tariff and not [y] will enable competitors to provide, and consumers to choose, tariffs including VoIP³. Similarly, if consumers decide that they are willing to pay more for completely unlimited/unmanaged Internet access then they will opt for those tariffs that make clear that P2P traffic is not throttled etc.

Nevertheless, it is not clear that additional regulation on transparency in respect of Internet services is required. While an adequate remedy, Everything Everywhere does not believe that there is a market failure as regards transparency etc that requires remedy. As set out above, Everything Everywhere considers that there are significant potential dangers in any mandated form of transparency and that this is best left to the market until such time as there is a clear case for regulatory intervention. More importantly however, there are significant incentives on competitors to increase transparency in order to create differentiated tariffs, enable more evolved traffic management, and differentiate themselves from each other. Hence, the combined factors of a competitive market and the need for traffic management will themselves drive transparency.

Question 6: Should the principles governing traffic management be the same for fixed and mobile networks?

Whether congestion can be relieved depends on a number of factors:

- the extent to which capacity can be expanded;
- the cost of additional capacity;
- the extent to which additional costs can be recovered, either through the retail market or through the wholesale market between ISPs and content/service or application providers; and.
- the extent of traffic growth.

In each case, there is a potential difference between the impact of these factors and the extent to which they can be addressed as between fixed and mobile networks.

Capacity expansion

The ability of ISPs to increase capacity is markedly different in the case of fixed and mobile networks.

In the case of the latter there are particularly significant constraints. Spectrum scarcity and the cost and difficulty of cell site installation mean that increasing the density and capacity of the radio access layer is problematic. Although individual sites can be augmented through (for example) additional sectorisation (dividing the cell in to further sections with additional antennae) this can only bring incremental improvements. Moreover, the cost and availability of additional backhaul capacity is a growing constraint. The pricing of high capacity backhaul is in many cases prohibitive (Openreach charges the same for second and subsequent Ethernet connections regardless of the incremental cost to itself) and even on those sites where this expense is unavoidable / can be recouped, the time taken for installation by the supplier is a significant hindrance. Additional regulation and or competition in backhaul will be determinative of the extent to which this particular bottleneck can be alleviated.

In fixed networks there are physical limitations to the amount of data that can be carried over an xDSL line. Although most consumers do not yet demand / have access to services that reach this threshold, the continued rate of growth makes this inevitable in the medium term, particularly given the growing consumption of television via the Internet. Hence, while in fixed networks congestion is less problematic, this will become an increasing factor.

This is particularly the case at the wholesale level however, where there are concerns around the BTW IP-Stream networks and their ability to increase capacity. With a current price of £122 per Mbps the cost of capacity is too high to ensure economic delivery over this infrastructure. Hence, even where ISPs and the home do not yet affront capacity constraints at present, this will become the case in the medium term, if only in the wholesale market.

³ As set out further above, there may be other parameters to such competition, such as guaranteed QoS, or "inclusive minutes" (e.g. VoIP not counting towards download limits etc).

Finally, at an upstream level it is unclear that peering will continue to provide a viable basis for traffic exchange. Few such links are in fact in balance and the growth of traffic means that this imbalance is growing among many peered partners. While it presently provides for a low cost means of traffic communication, these links are increasingly congested and there is little transparency downstream as to the extent of congestion or the QoS that can be expected by ISPs themselves. Capacity use could be significantly influenced through the incentivisation of greater efficiency at this level, although how/whether this can be developed will only become clear over time. The ability or otherwise of network operators to share the cost of capacity with OTT operators will therefore be a significant factor.

Cost of additional capacity

As will be clear from the above, the cost of expansion of capacity is different as between fixed and mobile ISPs, notably at the access layer.

However, both types of ISP rely on underlying fixed network infrastructure for backbone transmission. The continued high cost of backhaul/high capacity links means that capacity expansion at this level is increasingly expensive. This inevitably creates pressure to manage peak demand in order to control costs.

The extent to which additional costs can be recovered in the retail market

The market for both fixed and mobile internet access is highly competitive. Price is a significant parameter of consumer choice and accordingly ISPs are under continued competitive pressure to reduce costs. With a healthy number of competitors each attempting to expand market share there is continued competitive threat to any operator that attempts to increase prices.

The situation is inevitably slightly different as between fixed and mobile networks however, since the latter are increasingly congested already. Increasingly, mobile networks are introducing lower and more carefully supervised fair use thresholds, albeit at this stage, such policies are aimed only at constraining the most intense users of capacity. The thresholds introduced to date typically do not impact on the average customer for a given tariff and therefore serve only to significantly improve overall service by controlling capacity consumption by a minority of customers that use disproportionately high amounts of capacity.

The extent of traffic growth

Even with continued capacity expansion, traffic growth outstrips the ability of both fixed and mobile networks to maintain a lead on growing demand. Although the level of the bottleneck within networks differs, it is expected that all networks will become congested at peak times at some point over the development cycle (network expansion is inherently "lumpy").

The extent of traffic growth is particularly notable in the case of mobile networks, where the rapidly growing use of "smart phones" is driving increasing demand for capacity. While it may be expected that networks will introduce more stringent pricing and fair use policies, the ability to expand mobile network capacity is particularly constrained. The extent of traffic growth on mobile networks will therefore require mobile network operators to be particularly careful in their approach to traffic and demand management.

More generally, the fact that traffic will tend to grow to fill the available capacity must be borne in mind: traffic levels will inevitably continue to increase until there is congestion.

Question 7: What other forms of prioritisation are taking place? Do content and application providers also try to prioritise their services? If so, how – and how does this prioritisation affect other players in the value chain?

The role played by applications and content/service providers, as well as (in the case of mobile) device manufacturers, ought not be overlooked. ISPs do not control the speed/reliability etc of content or service delivery alone. This is also significantly determined

by the content and service providers themselves, who may or may not choose to host their content “near” the end user. For example, for high volume/low value services (e.g. video) Google chooses to route traffic via peering arrangements (with no guaranteed QoS) whereas for others it chooses to pay transit in order to guarantee the end user experience. Hence, they directly influence the QoS delivered to the end-user. More generally, the lack of any incentive on them to optimise the use of network capacity (most constrained among mobile networks) means that traffic management must in part be employed to make up for the lost opportunity to increase efficiency by content/service providers. Given the increasing volumes and imbalance of traffic there is a growing case for economic incentives for a more efficient use of network capacity to ensure the long-term sustainability of the open internet: commercial arrangements for the transport of IP traffic should better reflect the value of network capacity, encouraging a more efficient use of the network by the main sources of internet traffic, content and application providers.

Moreover, given the emphasis of dominance in the net neutrality debate, it is perhaps appropriate to consider the role of handset manufacturers, search and content providers and aggregators in this debate: while ISPs are a key element of the value chain, their role and influence is increasingly subsidiary to the global companies which drive the volume of traffic at issue, and which increasingly hold key consumer information, relationships and the content and services to which they seek access. Any net neutrality debate must therefore explicitly consider the roles of over the top (“OTT”) operators such as Google (You Tube, search/advertising, the Android operating system, etc) and Apple (iPhone, the Apple app Store, Apple TV etc).

It follows that:

- the speed/reliability etc of content/service delivery is a factor of two distinct but related markets, retail demand and content/service supply, where ISPs sit between two sets of potential customers with overlapping interests; and
- there are significant concentrations of market power elsewhere in the Internet supply chain which are relevant to the delivery and availability of competing services

Question 8: In the case of managed services, should the same quality of service conditions and parameters be available to all content/application/online service providers which are in the same situation? May exclusive agreements between network operators and content/application/online service providers create problems for achieving that objective?

The situation in the UK and Europe presents a much lower the danger of discriminatory abuse of dominance by vertically integrated ISP/content companies than is the case in markets such as the US. This is both because of different regulatory regimes on network competition and the relevant market participants. The danger of discrimination by ISPs between their own content and that of third parties is therefore much lower. Furthermore, with powerful national and European level regulatory oversight the ability of any dominant operator to abuse this position is significantly curtailed and would rapidly be identified. Network competition and regulation in the UK market ensure that an ISP favouring its own content cannot have the same consequences as are often touted in the US debate.

The more significant issue is therefore one of whether content/service providers should be free to enter in to contracts with ISPs for particular services. Everything Everywhere does not believe there is a basis for excluding this opportunity. For example, any general prohibition on agreeing exclusive content (for example) may significantly reduce opportunities for innovation and investment, since exclusivity is typically an important guarantee to both parties to ensure that the risk and investment that they commit to is not undermined – and may often be determinative of whether a new product is brought to market (subject of course to competition laws on foreclosure etc). While there are few relevant examples in the context of the Internet it is not difficult to imagine potential scenarios. For example:

- High definition online gaming.
HD online gaming requires expensive hardware that is rapidly out of date as technology and games evolve. Even where a user invests in this, their investment will lie idle for the majority of the time. Were a company to provide HD gaming via the Internet it could reduce the cost of hardware to users, since they would not need to meet the initial

investment cost and would be able to access processing power the cost of which could be recouped 24/24hrs, therefore making it a significantly more efficient proposition. However, HD online gaming where the core processing power is provided by a third party still requires that the output (namely the graphics) is communicated back to the user via the Internet, which requires high speed broadband and a significant level of guaranteed QoS: "best efforts" would simply not deliver a reliably usable service. It would therefore be attractive to an HD gaming company to be able to enter in to an agreement with an ISP for guaranteed QoS and market this to end users.

- Smart meters

Power utilities generate and distribute electricity in a gamble to predict where and in what quantities it will be consumed. The lack of any intelligence or communication in the distribution network prevents any interaction with users and appliances. A smart grid where consumption could be actively managed could enable far more efficient use of generation and distribution capacity (the cost of which is similarly driven by peak demand⁴) and create the possibility of a variety of differentiated tariffs for consumers⁵.

A smart grid will rely on communication between power generators and consumers, via smart meters and appliances. However, the amount of data required to be transmitted between utilities and smart meters/appliances is very small, since it only need communicate a small number of basic figures/commands. Consequently, high bandwidth is not a prerequisite.

It would therefore be inappropriate to prevent ISPs/utilities agreeing to the provision of services that have a set of QoS parameters that vary significantly from that which would ordinarily be required by a typical consumer/commercial customer. Minimum QoS requirements imposed under the auspices of net neutrality would therefore be to impose higher minimum costs on the roll out of smart networks.

It follows that there are a great deal of potential positive opportunities for innovation and differentiation through the exercise of contractual freedom among Internet actors, which would lead to direct consumer benefits. While it is clear that there are potential negative outcomes also, in a competitive market the scope for these will be extremely limited, as companies which attempt such approaches will be punished by consumers. Hence, so long as the market remains competitive, we do not believe that, on balance, there is any basis on which to preclude such arrangements being entered in to, or for requiring them to be made available universally.

Question 9: If the objective referred to in Question 8 is retained, are additional measures needed to achieve it? If so, should such measures have a voluntary nature (such as, for example, an industry code of conduct) or a regulatory one?

As should be clear from the above, we do not believe that it should be an objective of any party to require and/or impose uniformity. To attempt to achieve such an objective would ignore the right of the consumer to choose, and preclude the opportunity for companies to provide choice. It would significantly undermine innovation, service and the consumer benefit of a competitive market.

In our view, regulatory intervention to dictate the basis on which companies deal is only justifiable where there is market failure and a company is abusing its position to pursue its own interests, with the freedom to ignore consumers and competitors. At this point, regulatory intervention is of course necessary.

⁴ A 2006 study in the US found that a 5 percent drop in peak demand would enable the 625 infrequently used "peaking power plants" and associated power delivery infrastructure to be shut down. Since peak power plants (which are run only at peak times to meet short term demand spikes) have high variable costs and generate only a small proportion of their total potential output, such a reduction would yield an annual savings of \$3 billion, or a present value of \$35 billion over next two decades: Source – The Brattle Group: <http://sites.energetics.com/MADRI/pdfs/ArticleReport2441.pdf>

⁵ A smart grid could determine at what time particular appliances were used, according to demand levels/the prevailing price. Potential tariff structures include those where users obtain low prices in return for interruptible supply to high consumption appliances or applications (e.g. washer/dryers, car recharging)

However, the basis for such intervention already exists through the ability of NRAs to impose SMP conditions of non-discrimination etc. The legal basis for measures is therefore already in place and no further legal measures are required.

Question 10: Are the commercial arrangements that currently govern the provision of access to the internet adequate, in order to ensure that the internet remains open and that infrastructure investment is maintained? If not, how should they change?

Historically traffic levels were significantly lower and were more evenly spread. More recently however, notably with the explosion of online media content and services, traffic levels have increased dramatically. Moreover, this traffic/content is generated by a relatively small number of global enterprises.

It follows that historic arrangements which relied on a balance of traffic among a range of ISPs are coming under strain, as network operators struggle to fund the required expansion of capacity from the retail market alone. Not only does this reduce the amount of capacity that can be invested in (since consumers are highly price sensitive) but it is inefficient, since a two sided market is having its costs recovered from only one side.

The conclusion is not that content/service providers must be charged in future, but that the commercial arrangements currently in place will likely evolve. The extent of this evolution, and the scope of the traffic/links/peers etc to which it relates, cannot be foreseen as it will be dictated by the market as it develops. The important point however is that the current Internet and the commercial arrangements behind it have developed without regulation and there is no reason to conclude that the continued evolution of both must be regulated simply because this will create winners/losers. So long as no company is able to abuse a dominant position, the evolution of the Internet and its underpinnings will continue to develop in unpredictable and innovative ways, to the benefit of all.

Question 11: What instances could trigger intervention by national regulatory authorities in setting minimum quality of service requirements on an undertaking or undertakings providing public communications services?

Imposing a set of QoS requirements would be, in the opinion of Everything Everywhere, both inappropriate and dangerous. Not only is there no basis for imposing QoS (given the absence of dominance in the retail broadband market, evidenced consumer harm and the inherent difficulty of imposing QoS on a network of networks), but this would in fact result in a net detriment to consumer benefit, since it would reduce the potential for innovation and harm competition.

The former point (innovation, through the development of differentiated tariffs/services etc) has already been explored above. However, it is important not to overlook the potential harm to competition of imposing QoS requirements. Not only would it entirely prevent certain points of competition (e.g. markets for innovative products outside the precise minimum criteria required), but it would significantly reduce it on those parameters that were regulated. By in effect dictating a particular set of services, competitors would be drawn to compete only on these parameters, since these would inevitably be elevated in terms of their significance. Moreover, there is a significant danger that the particular tariff structure that evolved around the deemed set of QoS requirements would stagnate, evolving only with changes in the regulated QoS parameters around which they were framed.

It follows that we believe that the imposition of QoS can only be valid where there is market failure. This is not to say that transparency over QoS is a bad thing however, only that there is no basis for this to be regulated. In a competitive market transparency over QoS will continue to evolve and expand, notably as companies attempt to differentiate themselves from their peers, by developing a range of services/platforms.

Question 12: How should quality of service requirements be determined, and how could they be monitored?

There would be significant difficulty in policing compliance. This is because the actual level of service provided over a given connection depends on a wide range of factors. As noted above, the ISP with which a consumer/business contracts is only the final stage in the connection to the services/content to which that customer seeks access. The level of service provided via the ISP is contingent on the service provided by the ISP, but is not determined by it. For example, an ISP could have an entirely uncongested fibre network providing almost unconstrained speeds of access up to 1GBps, but nevertheless be unable to deliver HD video where the content provider's servers to which the customer is seeking access are themselves congested or faulty. Accordingly, any QoS requirement would be of only very limited relevance unless this was guaranteed on an end to end basis by all the actors in the supply chain.

Question 13: In the case where NRAs find it necessary to intervene to impose minimum quality of service requirements, what form should they take, and to what extent should there be co-operation between NRAs to arrive at a common approach?

Whilst in competitive markets net neutrality is primarily about transparency, in markets where regulators believe they may need to intervene they must look beyond the final layer of the internet at the point between ISPs and retail consumers and have regard to the market power that has developed elsewhere in the market. This is particularly the case in the event that any QoS remedies are considered. Not only do these have the potential to dampen competition (by removing competition on measured aspects of quality), but they may not be sufficiently targeted, as the ISP cannot itself determine all elements of the supply chain. More generally, the regulator must look beyond the role of the ISP and to the growing importance of the major (and relatively few) companies that are behind the demand for and delivery of content and applications to consumers.

As a consequence, while in competitive markets net neutrality is primarily about transparency, in markets where regulators believe they may need to intervene they must look beyond the final layer of the internet at the point between ISPs and retail consumers and have regard to the market power that has developed elsewhere in the market. This is particularly the case in the event that any QoS remedies are considered. Not only do these have the potential to dampen competition (by removing competition on measured aspects of quality), but they may not be sufficiently targeted, as the ISP cannot itself determine all elements of the supply chain. More generally, the regulator must look beyond the role of the ISP and to the growing importance of the major (and relatively few) companies that are behind the demand for and delivery of content and applications to consumers.

In summary, regulation of the Internet should not and cannot be driven through the ISP layer alone. Not only is this market competitive, but it is one of only a number of layers in the delivery of internet services. With the cost of capacity to the consumer determined by peak demand and the data intensity of particular services/applications, networks must balance network investment with costs in order to remain competitive

Question 14: What should transparency for consumers consist of? Should the standards currently applied be further improved?

Care should be taken to ensure that ISPs in the Internet supply chain are required to be transparent about what they are providing in a particular tariff/package, in order that consumers/businesses can ensure that they purchase the appropriate service for their requirements and in turn communicate demand to and competitive pressure on providers. However, to these same objectives should be shared by other actors in the supply chain, since they are equally involved in the delivery to the consumer and absent a full picture of what is going on, the value of information about one aspect of service is inherently undermined.

Everything Everywhere believes that net neutrality regulation, if any, should focus on ensuring that there is transparency in the market, albeit that we believe that the market will ensure this itself. Nevertheless care should be taken in the definition of any such transparency measures,

since they may come to define competition, and in doing so constrain it by precluding and/or devaluing innovation in services that do not fit the defined metrics.

For example, it may be superficially attractive to require all ISPs to provide a clear indication in their tariffing literature around a set of core metrics, by displaying these in any easy to understand table:

Price	£10
Maximum speed	10 Mbps
Download limit	10 GB
Traffic management	
Throttling P2P	Yes
Video prioritisation	Yes
VoIP	Yes

If such information was standardised and required as a minimum it would enable consumers to make quick informed decisions on the appropriateness of a given tariff, and to compare between different ISPs.

However, such high level standardisation would also present dangers. Most importantly, it would prevent ISPs competing on more nuanced offers that could not be pigeon-holed according to the defined transparency criteria.

For example, the above format would limit/preclude an ISP from the following aspects of differentiation

- Price: £10 for the first 50GBs of use per month, 50p per GB thereafter
- Video prioritisation: first 10GB prioritised, best efforts thereafter
- Video prioritisation: only on fixed broadband line, not on bundled mobile proposition
- Video prioritisation: guaranteed QoS for all [iTunes/Media Player/iPlayer] content between 5-10pm
- Throttling P2P: only between 5 and 10pm
- Throttling P2P: not for the first 10,000 customers
- VoIP: first 100 minutes do not count towards download limit

...and so on.

While none of the above may in fact be commercially viable etc, their purpose is to illustrate that there are potential variations around the standard parameters which may be attractive to certain consumers but which may be precluded by transparency criteria that effectively prevent their proper marketing. It is sufficient to look at the wide range of mobile tariffs offered to appreciate that consumers pursue a very wide range of services and combination of services/products and that limitations on the market in the form of transparency criteria which effectively limit the range of potential aspects of differentiation will reduce consumer choice, and therefore benefit.

Everything Everywhere therefore believes that any transparency requirements imposed should be carefully consulted on. It may be appropriate to require providers to publish clear information on x, y and z parameters but not dictate the form of the information provided for instance. This would ensure that consumers were provided with the relevant information, but would not preclude operators developing more nuanced tariffs that did not provide a binary choice on particular aspects of service.

Question 15: Besides the traffic management issues discussed above, are there any other concerns affecting freedom of expression, media pluralism and cultural diversity on the internet? If so, what further measures would be needed to safeguard those values?

Everything Everywhere does not believe that there are any particular issues that we have not already registered above. It would however emphasise that the Commission must look beyond the ISP layer and consider the wider "supply chain" of the Internet, since "the freedom of expression, media pluralism and cultural diversity of the internet" is in fact largely determined by the content and services made available over it. To the extent that ISPs broadly provide access to all online content/services, the Commission must consider the role

of the companies that generate and provide this in the character and development of the Internet.