

available at www.sciencedirect.comwww.compseconline.com/publications/prodclaw.htm

**Computer Law
&
Security Report**

Regulating the Internet

Net neutrality: A user's guide[☆]

Paul Ganley, Ben Allgrove

Baker & McKenzie LLP, London

ABSTRACT

Net neutrality is a complex issue that has generated intense levels of political discussion in the United States, but which has yet to attract significant attention from regulators in the UK. Nevertheless, the question of whether network operators should be prevented from blocking or prioritising certain network traffic or traffic from particular sources is a significant one for a wide range of stakeholders in the digital networked economy. Network operators contend that the build costs for the next generation of networks are so high that they must be permitted to monetise their control over this infrastructure as efficiently as possible. Meanwhile, an eclectic mix of interests including content and service providers, free speech and special interest groups and entertainers, argue that net neutrality regulation is necessary to guarantee that the Internet's core values and social utility are preserved. This article offers an introduction to net neutrality from a UK perspective. The authors explain the technical, commercial, political and legal considerations that underpin the issue and suggest that, whilst net neutrality regulation in its strongest incarnation is not practical or desirable, a level of regulatory action designed to enhance the choices of end users is the best way forward.

© 2006 Baker; McKenzie LLP. Published by Elsevier Ltd. All rights reserved.

1. Introduction

*Controlling who gets the fast lane is tantamount to giving control...*¹

In 1999 Cisco introduced a new type of router that enabled network operators ("operators") to inspect data packets flowing through their networks. The router allows operators to prioritise or de-prioritise certain packets of

data or even drop them from their network altogether. This technology, and its more advanced successors,² allow operators to choose how to handle data packets for commercial or policy reasons as opposed to the network performance reasons originally envisaged by Cisco. Packets can be favoured because they originate from a preferred source. Likewise packets can be de-prioritised or even blocked simply because they originate from a non-preferred source. This prioritisation or de-prioritisation of data packets is

[☆] The views expressed in this article are the authors' own and do not necessarily reflect the views of Baker & McKenzie LLP or its clients. All errors and omissions remain the sole responsibility of the authors.

¹ A Davidson, Google's Washington counsel, as quoted in R Ascierio, 'U.S. House Neuters Net Neutrality' *ComputerWire* (12 June 2006).

² A recent brochure from Cisco for an application called Cisco Service Control Engine describes a "deep packet-inspection engine" that allows network operators to "identify, classify, monitor and control traffic" through an "application-aware and subscriber-aware" system. Underneath the jargon lies an Internet policy that suggests discrimination based on the users and uses of the network. See 'Cisco Service Control: A Guide to Sustained Broadband Profitability' (Cisco White Paper, 2005) available from: <http://www.democraticmedia.org/PDFs/CiscoBroadbandProfit.pdf>.

0267-3649/\$ - see front matter © 2006 Baker; McKenzie LLP. Published by Elsevier Ltd. All rights reserved.

doi:10.1016/j.clsr.2006.09.005

often dubbed “access tiering” and it is at the core of the “net neutrality” debate.³

The ability to handle data on different network tiers has ignited a high-profile debate in the United States about whether or not operators should be allowed to discriminate between data packets and, therefore, whether regulatory intervention is needed to constrain how operators run their networks. This debate has prompted many to reconsider what public interest values are promoted by a “non-discriminatory” or “neutral” Internet and whether access tiering threatens that public interest. Importantly, the net neutrality debate is one which is now gaining traction in Europe. It is a debate which takes place in the context of various recent episodes that raise similar policy questions. Episodes such as Yahoo!’s dealings with the French courts on the question of the sale of Nazi memorabilia,⁴ Google’s forays into China⁵ and the debate about who should control ICANN.⁶ These episodes force governments, and society, to confront the question of how and whether the Internet should be regulated.

The net neutrality debate is often framed as having just two sides. On one side are the operators. In the US, the most vocal of these have been companies like AT&T, Verizon and Comcast. The operators argue that the increasing demands placed on the modern Internet require a level of investment that can and will only occur if the Internet is efficiently commercialised. They say that this commercialisation must involve the ability to implement a “user pays” model for the use of their networks and, hence, the Internet; those who make high use of and profit from the Internet, should, the operators say, pay for that use.

The other side of the debate is more complex and is characterised by an eclectic coalition of content and service providers, such as Google, Intel, Yahoo!, eBay and Amazon, anti-regulation advocates, entertainers, like REM and Moby, free speech groups, like Free Press, and others such as the Christian Coalition, National Religious Broadcasters and the Gun Owners of America. The message that these groups and individuals send out is that access tiering threatens the core values and social utility of the Internet and that governments must intervene to prevent access tiering from occurring.

³ The net neutrality debate is often characterised by the use of emotive terms such as “discrimination”, “neutrality”, “freedom” and “democracy”. The use of such rhetoric often clouds consideration of the issues at play because it elides what are often quite nuanced and diverse issues into bi-polar “pro” and “con” camps. Further inflated claims foretelling the end of the Internet, the stagnation of broadband deployment and the death of free speech often characterise this debate. The authors prefer, where possible, to use the technological term “access tiering” because it is objective, allowing the pros and cons of the technological ability to prioritise and de-prioritise data packets to be assessed without falling into the trap of over-simplification.

⁴ *L.I.C.R.A and U.E.J.F v. Yahoo Inc. and Yahoo France*, Interim Court Order, Paris County Court (20 November 2004).

⁵ See Clive Thompson, ‘Google’s China Problem (And China’s Google Problem)’ *New York Times* (23 April 2006) available from: <http://www.nytimes.com/2006/04/23/magazine/23google.html?ex=1303444800&en=972002761056363f&ei=5090>.

⁶ See ‘Who owns the web?’ *The Sunday Herald* (20 November 2005), p. 24.

In this article, the authors tackle the net neutrality debate. In doing so, they show that net neutrality is not simple and bi-polar. Rather, it is a complex and fascinating issue that must meld the public interest with legal, practical and commercial considerations. At the end of the day, there is no inherently “correct” position. Compromises must be sought and reached. These compromises must balance the increasing demand for investment which the modern Internet occasions with the genuine concerns that the sourcing of that investment should not undermine the largely unfettered exchange of information that has characterised the development of the Internet to date and which, in the minds of many, is what has made the Internet such a powerful social force in such a short space of time.

2. The technology

2.1. Regulating the Internet: the “layers principle”

Before turning to discuss the net neutrality debate it is necessary to consider to what extent the Internet is currently, and can be further, regulated.

The idea that the Internet should be unregulated, and indeed could not be regulated, reached its zenith with the publication of John Perry Barlow’s ‘Declaration of the Independence of Cyberspace’ in 1996.⁷ Barlow’s conception of the Internet as an independent ‘space’ or institution has taken root with those who argue that the Internet should be “free” – both of government regulation and commercial distortion. As the Internet has developed, however, it has become increasingly clear that it is subject to many of the same regulatory forces as other social institutions and fora.⁸ That the Internet *can* be regulated and influenced (both by governments and those that operate its infrastructure) is no longer in issue; the question now being asked is, should it be regulated, and if so, how?

In order to understand how the Internet can be regulated or influenced, one must understand the various “layers” of Internet topology and how each of these layers is susceptible to regulatory pressure.

Broadly speaking, the Internet is comprised of three layers: the physical layer, the logical layer and the content layer.

- The *content* layer is made up of the content, information and other meaningful statements that individuals using the Internet perceive, act on, laugh at and share.
- The *logical* layer describes the series of algorithms and standards – including TCP/IP, HTTP and HTML – that allow content layer materials to be understood and transmitted in machine readable form; it is one part of the “machinery” of the Internet.
- The other part of that machinery is the *physical* layer, which includes the tangible objects – computers, wireless devices,

⁷ John Perry Barlow, ‘A Declaration of the Independence of Cyberspace’ (8 February 1996) available from: <http://homes.eff.org/~barlow/Declaration-Final.html> (“you weary giants of flesh and steel...you have no sovereignty where we gather”).

⁸ See, e.g. Jack Goldsmith and Tim Wu, *Who Controls the Internet: Illusions of a Borderless World* (Oxford University Press 2006).

wires, routers etc. – that connect individuals to the Internet and one another.

At each of these layers we witness controversy, influence and regulation of varying kinds.

At the content layer, these battles have primarily been fought in the realm of the private enforcement of copyright and other intellectual property (“IP”) rights. The result of these private enforcement battles has generally been governments strengthening laws to offer a narrower set of permitted uses of protected content and the increased criminalisation of infringement of IP rights. Operators, meanwhile, have successfully resisted attempts to make them function as “rights police” so long as they take reasonable measures to prevent infringement and act when infringement is brought to their attention.⁹ Other battles at the content layer have been fought around database rights, pornography, gambling and defamation, to name but a few. The battles at the content layer tend to be about translating physical world controls into the digital medium.

At the logical layer, these battles tend to focus on the technology which underlies the Internet. Examples include actions against peer-to-peer software providers¹⁰; the controversy surrounding the implementation of anti-circumvention laws in the US and the Europe; and the continual efforts to ensure that the domain name system functions effectively.¹¹ Another example is the Free/Open Source Software (F/OSS) movement that has been embraced by many software vendors and which favours open, but not necessarily free, access to some of the key components of the logical layer.¹²

At the physical layer, increasingly fractious arguments characterise the debate, including arguments about areas as diverse as free or low cost municipal wireless Internet (“wifi”) systems; the regulation of hardware – e.g. personal computers and other devices that process content – through the imposition of “trusted system” architecture, such as the “broadcast flag” in the United States; and the government’s ability to intercept communications and to control encryption technologies.

Understanding these layers and the battles at play at each of them, makes it instantly apparent that Internet regulation is about more than just law. Legal, technical, social and market based rules and norms interact to determine the dynamics of the Internet. The net neutrality debate is no different. It is a debate about regulation and influence at the interface of the logical and physical layers. Before we delve into the details of the controversy, however, it is worth understanding what the physical Internet first looked like and how it continues to look, largely, to this day.

⁹ Directive 2000/31/EC on Certain Legal Aspects of Information Society Services, in Particular Electronic Commerce, in the Internal Market, articles 12–15; Digital Millennium Copyright Act, §512.

¹⁰ See *A&M Records, Inc. v. Napster, Inc.* 239 F.3d 1004 (9th Cir. 2001); *A&M Records, Inc. v. Napster, Inc.*, 284 F.3d 1091, 1096 (9th Cir. 2002); *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster, Ltd.*, 125 S. Ct. 2764; and *Universal Music Australia Pty Ltd v. Sharman License Holdings Ltd* [2005] FCA 1242.

¹¹ See, e.g. Milton Mueller, *Ruling the Root: Internet Governance and the Taming of Cyberspace* (MIT Press 2004).

¹² See, e.g. Steven Weber, *The Success of Open Source* (Harvard University Press 2004).

2.2. The end-to-end principle

The Internet was designed as a “dumb” network. Its central function – implemented via the TCP/IP protocols – is to pass packets of data, via “pipes”, along a chain of “nodes” until they reach their destination. The nodes do not ask questions about the sender of the packet, the recipient, or its content; they simply receive them, analyse the address information and pass them on to the next node. This dumb network treats all packets equally – a principle referred to as “bit parity” and often encapsulated in the phrase “end-to-end” design.¹³ In a dumb network, intelligence is incorporated in the applications that sit at its edges, or “ends”. The network itself is not designed with any particular applications in mind¹⁴; it will work with any application that someone cares to invent. The applications at the ends may themselves perform “intelligent” functions – like blocking junk emails, suppressing pop-up ads, identifying possible viruses, or flooding P2P networks with spoof files – but the core of the Internet’s infrastructure could not care less.

Many argue that the end-to-end principle has been instrumental to the remarkable growth and success of the Internet by increasing competition and innovation between content and service providers and by enabling a wider variety of applications to connect to and use the Internet. For instance, it has been suggested that “had the original Internet design been optimized for telephone-style virtual circuits...it would not have enabled the experimentation that led to the protocols that could support the World Wide Web”.¹⁵ The net neutrality debate is really about whether the Internet should retain its end-to-end design or whether the operators, who own and control various aspects of the physical layer, should be permitted to “discriminate” amongst the data that passes across their networks by access tiering.

2.3. Preference without access tiering

...the Internet already has slow and fast lanes. Companies such as eBay and Google plug into it through big pipes and store their data on servers around the world so that their pages load more rapidly. Telecoms companies exchange traffic in order to make it travel faster than it would through the public hubs that were the foundation of the Internet.¹⁶

It is important to appreciate that the notion that we currently have a “neutral” Internet is simply false. Even before we get to access tiering, there is a range of existing strategies that one can pursue to “prioritise” data over other network

¹³ J.H. Saltzer et al., ‘End-to-End Arguments in System Design’ 2 *ACM Transactions in Computer Systems* (November 1984) pp. 277–278 also available from: <http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.txt>.

¹⁴ This is in contrast to other networks, such as a circuit-switching telephone network.

¹⁵ David Reed et al., ‘Comment on Active Networking and End-to-End Arguments’ 12 *IEEE Network* (May/June 1998) pp. 69–71 also available from: <http://web.mit.edu/Saltzer/www/publications/endtoend/ANE2ecomment.html>.

¹⁶ J. Gapper, ‘Is it too soon to impose net neutrality’ *Financial Times* (1 May 2006) 15.

traffic. For example, an IP transit arrangement between a service provider and an operator will typically include service level guarantees, which give the operator a commercial interest in ensuring that those service levels are met. Similarly, a service provider may pay an operator directly to host content, thus generally guaranteeing a higher quality of service and reliability. Other strategies include utilising intermediary service providers, such as Akamai, who purchase excess bandwidth from local access providers to host content in local caches at various locations around the globe so as to ensure that data requested by customers never has too far to travel.¹⁷ All these existing strategies, however, lie towards the “edge” of the Internet. They flow off the back of what users of the Internet are prepared to invest in their infrastructure and services.

2.4. Access tiering

In contrast to user led initiatives mentioned above, the access tiering models now advocated by some operators (and which are the subject of the net neutrality debate) are a more sweeping attempt to adjust the Internet’s default settings by placing control of the network in operators’ hands and allowing them to set the price for access.

- The “best efforts” rule, the existing default, treats all data packets the same; it is a first in/first out arrangement irrespective of origin or destination.¹⁸ It applies bit parity.
- “Needs-based discrimination”, on the other hand, treats packets according to the best efforts rule until such time as there is network congestion. At this point, certain packets – latency-sensitive ones for example – are prioritised and move to the front of the queue.
- “Active discrimination” occurs when operators inspect all packets and prioritise them in accordance with pre-defined rules irrespective of whether their network is congested. Needs-based and active discrimination overturn the best efforts rule and bit parity.

In documents filed with the Federal Communications Commission (FCC) in the US, for example, operator Verizon explained how it plans to use one laser in its fibre-optic cables exclusively for Verizon’s own video services. The other laser will be used to provide VoIP, IPTV and other Internet services for third parties. Experts who have analysed Verizon’s plans suggest that upwards of 80% of its network capacity has been earmarked for Verizon’s own services, leaving just 20% for other Internet users who may need their data to travel along Verizon’s pipes in order to get to their intended recipients.¹⁹

¹⁷ Whilst many claim the Internet is borderless, the physical location of data is in fact a critical determinant in how quickly it reaches end users. See, e.g. Goldsmith and Wu, *supra* n8 at pp. 49–63.

¹⁸ In times of congestion packets may be dropped, but not according to the properties of the individual packets themselves.

¹⁹ Catherine Yang, ‘Is Verizon a Network Hog?’ *BusinessWeek online* (2 February 2006) available from: http://www.businessweek.com/technology/content/feb2006/tc20060202_061809.htm.

Two further hypothetical examples illustrate the point.

- *Service provider discrimination*: an operator, such as BT, could enter into an agreement with Search Engine A under which A’s content is favoured over the content of its rivals including Search Engine B. At times,²⁰ this preference may be noticeable to the end user, possibly to the extent that users who become frustrated with Search Engine B migrate to Search Engine A because of its better performance.
- *Application discrimination*: alternatively, an operator could distinguish between applications, rather than providers. For example, the operator may decide that latency-sensitive applications such as VoIP or video streaming should be prioritised over less time sensitive packets such as those that make up emails or downloads.

Application and service provider discrimination can operate in tandem. So, for example, an operator who runs a telephony network, such as BT, could decide that VoIP competes with its core business of voice calls over circuit-switching networks. It may decide that given this competitive threat, it will de-prioritise all VoIP services. VoIP may nevertheless continue to cannibalise the operator’s existing revenue streams, in which case it may decide to enter the VoIP market itself and prioritise traffic originating from its own VoIP service above all other network traffic, thereby securing a competitive advantage for itself in two markets.

The commercial interest for operators in being able to access tier is obvious. So too is the reason for the fears that those opposed to access tiering express. Competitors fear being unable to compete. Non-commercial entities fear being unable to pay. Those that upset the apple cart fear being excluded altogether. With control of the Internet comes power; access tiering is potentially a very powerful form of control.

3. Recent US experience

We in Europe can get a flavour for the interests at play in the net neutrality debate by looking at recent US experience. The first thing of note is that for an issue that has received such a high level of political attention, there have been precious few examples of operators actually engaging in access tiering to date. The examples that do exist are noteworthy as much for their rarity as for the behaviour of the operator in question. The examples also concern *blocking*, rather than access tiering, and include the following:

- In 2004, North Carolina ISP Madison River blocked their DSL customers from using rival VoIP services. The FCC in the United States, acting on a complaint from Vonage, soon reached an agreement with the ISP requiring it to stop blocking VoIP calls and make a “voluntary payment” of \$15,000.²¹

²⁰ Particularly at peak times of network congestion or where the user is wishing to access large files.

²¹ Declan McCullagh, ‘Telco agrees to stop blocking VoIP calls’ *News.Com* (3 March 2005) available from: http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-05-543A2.pdf.

- Last year, Canadian telephone operator Telus blocked customers from visiting a web site sympathetic to the Telecommunications Workers Union (TWU) at a time when Telus and the TWU were involved in a labor dispute.
- Most recently, AOL has blocked emails that mention “www.dearaol.com”, an advocacy campaign opposing AOL’s pay-to-send email scheme.

The US debate has generally been about what operators *could* or *might* do unless net neutrality principles are enshrined in law, rather than what they are *actually* doing at the moment.

The second thing of interest that comes out of the US experience is that instances of overt blocking of content are limited and a consensus seems to be emerging amongst operators that blocking applications, content and specific service providers is not in their interests. AT&T Chairman Ed Whitacre for instance has pointed out that “Any provider that blocks access to content is inviting customers to find another provider. And that’s just bad business.”²² Similarly, James Speta has argued that the economics of telecommunications are such that “It is against the platform owner’s interest to attempt to monopolize content – even if the platform owner is a monopolist in transmission service”.²³ Speta posits that because customers derive value from interconnection *and* access to content “indirect network externalities” play a role in ensuring that network owners have the incentive to provide the most efficient set of content possible, which inevitably includes content from a wide range of providers.

The current strategy of US operators appears to be to accept that basic broadband speeds – say in the range of 1–2 MBs – should be generally available to all, but that incremental improvements in bandwidth over and above this basic level of “best efforts” service should be reserved for the services of the operators themselves or their *paying* partners.

Spurred on by the FCC’s active involvement in the Madison River case, many US legislative measures have been introduced aimed at enshrining net neutrality requirements in law. Various bills with self-explanatory names such as the “Communications, Consumer’s Choice and Broadband Deployment Act” and the “Internet Freedom and Non-discrimination Act” have been introduced, but, at the time of writing, all have stalled at various points in the legislative process.

A high-profile vote in the US House of Representatives in early June 2006 saw an attempt to enshrine net neutrality requirements in federal law fail. Under the proposed law “broadband network providers” would have had a general duty²⁴:

- Not to “block, impair, degrade, discriminate against or interfere” with a person’s ability to access, use or offer lawful content, applications or services over the Internet.
- To operate their networks in a “non-discriminatory manner” so that anyone is able to offer content and applications through the network without charge and with “equivalent or better” capability than the provider itself.
- To only prioritise *all* data of a particular type regardless of origin if enhanced quality of service is to be extended to certain data types.

The vote received a great deal of attention. eBay CEO Meg Whitman emailed more than a million eBay members urging them to press their representatives in Congress to support the legislation.²⁵ Likewise Google CEO Eric Schmidt wrote an open letter to Google users urging them to “take action to protect Internet freedom”.²⁶ On the other side it has been reported that telecom and cable companies in the US have been spending \$1 million per week on advertisements that oppose net neutrality legislation.²⁷

4. For and against network neutrality

4.1. The participants in the debate

Apart from the clear commercial interests of the operators and service providers themselves, the other interests in the net neutrality debate are as diverse as the activities that take place on the Internet.

- Hollywood studios hope to guarantee the quick and secure delivery of their premium content to paying customers, and see further revenue potential in the monitored use of their products.
- Content and service providers like Amazon, Google, Yahoo! and eBay want to ensure that another cost does not appear on their books and that they are not excluded from or discriminated against on the key parts of the Internet infrastructure they use to access their customers.
- Law enforcement agencies desire expanded capabilities to monitor and extract information from the sea of electronic data. Deep packet-inspection aligns nicely with this surveillance function.
- Hundreds of individuals, non-profit groups and businesses, from across the political spectrum, have come together as the ‘SavetheInternet.com’ coalition to promote net neutrality, which they refer to as “the Internet’s First Amendment”²⁸

²² Marguerite Reardon, ‘AT&T chief, FCC chair clarify on Net neutrality’ ZDNet (21 March 2006) available from: http://news.zdnet.com/2100-1035_22-6052239.html.

²³ James Speta, ‘Handicapping the Race for the Last Mile? A Critique of Open-Access Rules for Broadband Platforms’ 17 *Yale Journal of Regulation* (2000) 39, at p. 76.

²⁴ A copy of the proposed amendment to the COPE Act is available from: http://www.rules.house.gov/109_2nd/specialrules/2nd109/hr5252/109hr5252_markey20.pdf.

²⁵ Declan McCullagh, ‘eBay tries email in Net neutrality fight’ *News.Com* (1 June 2006) available from: http://news.com.com/eBay+tries+e-mail+in+Net+neutrality+fight/2100-1028_3-6079291.html?tag=nl.

²⁶ Eric Schmidt, ‘A Note to Google Users on Net Neutrality’ available from: <http://www.google.com/help/netneutrality.html>.

²⁷ Lawrence Lessig, ‘Network Neutrality: Critical Push’ (Blog post, 31 May 2006) available from: <http://www.lessig.org/blog/archives/003421.shtml>.

²⁸ ‘SavetheInternet.com Coalition Statement of Principles’ available from: <http://www.savetheInternet.com/=principles>.

To get a flavour for the fractious nature of this debate one need only to compare www.savetheinternet.com (pro-net neutrality regulation) and www.handsoff.org (anti-net neutrality regulation). The existence of these diverse interests creates the backdrop for the net neutrality debate and produces some delicious ironies – such as advocates of a “free” Internet calling for government intervention to ensure that it stays that way.

4.2. Investment in network infrastructure

The primary argument put by operators is that the pressures being placed on the Internet in terms of the number of users and the types of use are expanding rapidly. New applications – such as streaming video and voice telephony – are emerging. Furthermore, the Internet is increasingly being used in critical applications such as health monitoring and home security. These applications either sap bandwidth or demand high levels of service quality, both of which place extra burdens on the infrastructure built and maintained by the operators. The pipes that facilitate this global data exchange are beginning to buckle under the weight of their own success.

In the US there is talk of a “copper ceiling” being hit by 2007.²⁹ This will restrain Internet-based services to the ones currently offered where fibre-optic networking cannot be delivered. The cost of refitting all US copper with fibre optics is estimated at \$45 billion – a very significant investment that, regardless of one’s attitude towards net neutrality, has to be met somehow.

Similarly, there is increasing demand for security measures – designed to weed out spam email and malicious viruses – implemented at the “core” of the Internet, rather than at the user-interface. In short, operators want to give the dumb network an education. In doing so, operators are hoping that they will be able to capture a greater share of the “value” generated by the Internet.

Operators are now rolling out the next generation of networks to meet the increased demands of the modern Internet.³⁰ These networks come at an enormous cost. In 2006 AT&T expects to spend approximately \$20 billion dollars on its new Project Lightspeed. The amount spent by Verizon will be comparable. In the UK, BT is spending £10 billion in the next few years on its 21st Century Network (21CN) which involves the transition from a public switched telephone network to an IP-based packet switched network, including the provision of fibre-optic cables (instead of copper lines) to

²⁹ L. Turner, ‘Operators, vendors unite to drive IP services’ *Total Telecom* (24 May 2006).

³⁰ AT&T for example is in the process of constructing a high-capacity fibre-optic network dubbed “Project Lightspeed”. This network should serve 18 million US households by 2008 and will offer download speeds of between 20 and 30 Mbps; a 10- to 20-fold increase from today’s norm. To put this in context, whereas today’s DSL lines struggle to carry one standard-quality video stream, in tests AT&T’s new network could handle three standard-quality streams, one high-quality stream, Internet access and VoIP services simultaneously. See ‘Successful Technical Field Trial of IP-Based Video, High-Speed Internet Access’ AT&T Press Release (3 November 2005) available from: <http://att.sbc.com/gen/press-room?pid=4800&cdvn=news&newsarticleid=21874&phase=check>.

some areas. Telecom Italia estimates that over the next 3–4 four years, European operators will invest around \$97 billion in next-generation networks.³¹ These companies can expect to recoup some of this cost from end user customers via Internet access charges, but they would also like to extract money from content and service providers. AT&T Chairman Ed Whitacre, for instance, has vociferously complained that content providers get a free ride: “They don’t have any fibre out there. They don’t have any wires...They use my lines for free – and that’s bull. For a Google or a Yahoo or a Vonage or anybody to expect to use these pipes for free is nuts!”³² Likewise Verizon senior vice president and general counsel John Thorne has described service providers such as Google as “enjoying a free lunch that should, by any rational account, be the lunch of the facilities providers.”³³

In Europe, this point of view has been most strongly supported by Kai Ewe Rieke, the chief executive of Deutsche Telekom, who recently said “It shouldn’t be the case that infrastructure providers, like Deutsche Telekom, are always making the investments, while others profit of the back of those...these (web-based) companies should realise that it is us who will in the future guarantee network quality for their new applications”.³⁴

Operators argue that they should be allowed to fully exploit their property interests by charging certain content or service providers for enhancing the end-user’s experience. At a micro level, they argue, this will meet the needs of consumers and defray costs that would otherwise be passed on to them, whilst at a macro level it will hasten the deployment of the next generation of networks. Any problems can be addressed by anti-trust or competition law regulations, and absent a clear demonstration of market failure, market freedom should be endorsed.³⁵

The counter-argument runs along two lines. First, end users, content providers and service providers have for years been paying for network enhancements through subscription charges and bandwidth charges. They will continue to do so. Google, for example, has refused to share costs with operators claiming that “customers are already paying to support broadband access to the Internet...[and they] should have the freedom to use this connection without limitations”.³⁶ Likewise, even if there is no objection to the notion of the likes of Google

³¹ M. Odell and R. Wates, ‘A two-speed internet? Why network operators are flexing their muscles’ *Financial Times* (20 May 2006) 15.

³² See Christopher Stern, ‘The Coming Tug of War Over the Internet’ *Washington Post* (22 January 2006).

³³ See Arshad Mohammed, ‘Verizon Executive Calls For End to Google’s ‘Free Lunch’ *Washington Post* (7 February 2006).

³⁴ Odell and Wates, *supra* n31.

³⁵ These arguments are also supported by some in the legal and economic literature. See, e.g. Adam Theiner, ‘Are Dumb Pipe Mandates Smart Public Policy? Vertical Integration, Net Neutrality and the Network Layers Model’ 3 *Journal of Telecommunications and High Tech Law* (2005) 275 and Christopher Yoo, ‘Network Neutrality and the Economics of Congestion’ (forthcoming 2006) available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=825669.

³⁶ Preston Gralla, ‘Google: We Won’t Pay Broadband Cyberextortion’ (Blog post, 18 January 2006) available from: http://www.networkingpipeline.com/blog/archives/2006/01/google_we_wont.html.

paying for the bandwidth it uses, it should not be forced to pay additional rates determined by a third party based solely on the type of application it wishes to channel through the network.

The second counter-argument asserts that property rights rhetoric should not disguise the fact that operators are seeking to claim a share of the value created by others. Susan Crawford, for instance, asserts that the “powerfully romantic vision” of the “network builder” is being used to support claims for a share of future revenues generated by innovations at the content or application layer.³⁷ Analogies abound: road builders demanding commissions from travelling salesmen, Microsoft claiming royalties from the sale of works created using Word, or the booksellers’ revenue-share arguments in the ongoing dispute over Google Book Search.³⁸

In stark economic terms, this counter-argument posits that attempts by operators to generate revenue from the externalities that the network infrastructure helps produce should be restricted. So long as operators have sufficient incentives, in terms of economic returns, to build the next generation of network architecture, that is enough.

4.3. Innovation

The other side of the property rights coin is the question of innovation. Net neutrality enthusiasts maintain that access tiering will jeopardise the future of innovation online. They suggest that the end-to-end principle catalysed the intense levels of innovation that the Internet has spawned. The WWW, P2P software, VoIP, blogging tools and like innovations may not have emerged from a network infrastructure capable of discriminating between data packets.

Simply put, this argument maintains that one of the central virtues of the Internet is its ability to level the playing field for application and content development and support an environment where those start-ups or small providers with the most promising innovations – yesterday’s Google – could effectively compete against the monoliths of the day. At a recent US Senate hearing, Stanford law professor Lawrence Lessig argued that it is precisely this “competition on the merits” that is threatened when operators have the capability (to match their rational incentives) to discriminate amongst content or services.³⁹ If legacy business models and affiliated service providers are favoured, or if smaller innovators are unable to pay for prioritised access, the “creative destruction” that has characterised much of the last 10 years will be

³⁷ Susan Crawford, ‘Network Rules’ (working paper 2006) pp. 3, 19 and 20 available from: http://scrawford.net/courses/network_rules_for_22m.pdf. See also Susan Crawford, ‘The Possibility of Future Profits’ (Blog post, 28 January 2006) available from: http://scrawford.blogware.com/blog/_archives/2006/1/28/1727859.html.

³⁸ For more on this point generally and in the context of the ongoing dispute over Google Book Search see Paul Ganley, ‘Google Book Search: Fair Use, Fair Dealing and the Case for Intermediary Copying’ (working paper 2006) available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=875384.

³⁹ Testimony of Lawrence Lessig at Senate Committee on Commerce, Science and Transportation Hearing on “Network Neutrality” (7 February 2006) p. 9 available from: <http://commerce.senate.gov/pdf/lessig-020706.pdf>.

undermined and innovators will be dissuaded from investing in areas where network operators have a controlling hand over its deployment. With access tiering, Lessig warns, “innovators [will have to] now include in their calculation of risk the threat that the network owner might...tax a particular application. That increased risk will reduce application investment”.⁴⁰

4.4. Users rather than consumers

Those in favour of network neutrality suggest that those who support access tiering wish to recreate the offline broadcast-model of content distribution online – pre-packaged content fed to passive consumers.⁴¹ In truth, the Internet represents a far richer and deep-rooted set of content, applications and social connections. The argument, presented most forcefully by scholars such as Susan Crawford and Yochai Benkler, is that the Internet is “special”; a place which is user-centric and where opportunities for participation in content creation, dissemination and reuse offer significant social benefits along a multitude of paths tied to notions of free speech, individual autonomy and democratic participation.⁴² Understanding this argument is important as it enlarges the sphere of influence of the debate and suggests that competition in the broadband market may not of itself be enough to fully protect the public interest.

The argument posits that the Internet has spawned a vast array of web sites, applications and basic resources that tap into the production capacity of individuals in the digital networked space. These sites, applications and resources place end users at the heart of the “informational universe”. Their existence supports the notion that the Internet is a fundamentally different resource from that with which we are familiar offline: a frighteningly efficient communications mechanism and global marketplace for sure, but also an information repository of incredible breadth and diversity and a huge talent pool capable of being harnessed and managed across a range of creative projects. True, many of these individual-led initiatives will continue to function perfectly well if they are forced to operate through lower tiers of the Internet backbone, but this is simply a function of the Internet’s text-based origins. The richer multi-media environment we are constructing means that some existing resources – online video repositories “youTube.com” or “YouAre.tv” for instance – and many future initiatives may be stifled or never see the light of day if access tiering is allowed. The proponents of net neutrality regulation who make this argument are basically asking for limitations to be placed on operators’ private property interests for the sake of the public interest that exists in the maintenance and development of this “public space” provided by the Internet.

The problem, of course, is that damage to these public interests is difficult to measure or even predict. As Frischmann

⁴⁰ *Ibid* at p. 6.

⁴¹ See, e.g. Yochai Benkler, ‘From Consumers to Users: Shifting the Deeper Structures of Regulation towards Sustainable Commons and User Access’ 52 *Federal Communications Law Journal* 561 (2000).

⁴² Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (Yale University Press 2006); Crawford, *supra* n37.

puts it in a recent article analysing the optimal way for managing infrastructure resources that serve as “inputs” for a wide range of downstream productive “outputs”:

*The social costs of restricting access to public or social infrastructure can be significant and yet these costs evade observation or consideration within conventional economic analysis. Initially, we may analyze the issue as one of high transaction costs and imperfect information. Yet, even with perfect information and low or no transaction costs with respect to input suppliers and input buyers, input buyers would still not accurately represent social demand because it is the benefits generated by the relevant outputs that escape observation and appropriation.*⁴³

These effects will be particularly severe in circumstances where operators differentiate between download and upload speeds as most companies currently do.⁴⁴

The desire of operators to follow the broadcast-model is understandable in the context of companies versed in the industrial age economics of mass-media and facing large network build costs. But it is the Internet’s divergence from this model to a user-centric “read/write”⁴⁵ model – with its own set of attendant problems relating to information glut and viable accreditation – that leads many to argue that preventing access tiering is essential to preserving and nurturing the public interest in the Internet.

4.5. Resisting the regulatory impulse

One of the surprising aspects of the whole network neutrality debate – and a fascinating irony – is that those who typically oppose “Internet regulation” are broadly in favour of legislation designed to safeguard net neutrality. However, the answer to this apparent contradiction is obvious if one adopts a broad concept of regulation. Far better to have a freely debated, transparent and ubiquitous public regulation, argue the proponents of net neutrality, than a series of commercial and opaque private relationships determining the nature of the Internet. Both outcomes “regulate” the net in the sense that they constrain or promote particular types of behaviour, so it’s less a question of *if* we regulate, than *how* we regulate.

Opponents of net neutrality regulation emphasise the dangers of public regulation of the Internet. For example, Republican Senator John Sununu recently said “a heavy regulatory hand kills incentives to develop new products, deploy new technologies and that ultimately will be something

consumers will feel and respond to in a very negative way.”⁴⁶ In one sense their argument is perfectly sound: we should not confuse the unregulated Internet with the highly regulated telecoms sector that has, for many years, been subject to common carrier regulations in many countries including the US and the UK. Net neutrality regulations, they argue, are still regulations, which fail to recognise the early explosive growth of the Internet in a sphere largely free from government interference. This regulation will inevitably lead to unforeseeable and unintended consequences.

This argument, however, is open to two basic criticisms. First, it may be more difficult and expensive to regulate when a position has become entrenched. Or as Lemley and Lessig put it: “To say there is no reason to use a seatbelt because there is always the care of an emergency room is to miss the extraordinary cost of any *ex post* remedy”.⁴⁷ Second, a regulator may have difficulty seeing and certainly measuring the real world effects of a problem when those effects are basically something – an innovation say – that has not happened.

5. A transient debate?

The net neutrality debate is in one sense a network *engineering* debate. Operators tell us that a tiered Internet – fast, slow, and any number of iterations in between – will improve network efficiency. The more managed it is, they argue, the greater efficiencies there will be in traffic handling. Others are not so sure, claiming that introducing intelligence necessarily introduces complexity that can actually impede network performance.

At a recent US Senate hearing on net neutrality, Gary Bachula of Internet2, a non-profit project to build a super fast and extensive network, testified that whereas the Internet2 engineers had originally assumed that they needed to find technical ways to prioritise certain kinds of traffic such as streaming video or video conferencing, “all of our research and practical experience supported the conclusion that it was far more effective to simply provide more bandwidth”.⁴⁸ This comment captures the attitude of some who claim that the net neutrality debate is necessarily transient in nature as all our bandwidth needs will eventually be met, alleviating the need for any form of tiering. Indeed, in countries with faster broadband access – South Korea or Japan, say – the question of net neutrality has hardly arisen. Once the pipes are there, there is no incentive for operators to access tier, as their consumers want access to the entire Internet. New or alternative broadband technologies such as BPL (broadband over power lines), broadband via satellite and WiMax may also render the debate moot.

⁴³ Brett Frischmann, ‘An Economic Theory of Infrastructure and Commons Management’ 89 *Minnesota Law Review* (2005) 917, at p. 976.

⁴⁴ See, e.g. Bulldog Broadband ‘Broadband Options’ available from: <http://www.bulldogbroadband.com/residential/products/broadband/#> (up to 16 Mb/s download speed compared to 1 Mb/s upload speed). BT’s download speeds reach 8 Mb/s, whereas typical upload speeds are in the range 64–832 Kb/s (email correspondence with BT, on file with the authors).

⁴⁵ See Mark Lawson interview with Tim Berners-Lee on Newsnight, 9 August 2005. A transcript of the interview is available from: <http://news.bbc.co.uk/1/hi/technology/4132752.stm>. See generally Tim Berners-Lee, *Weaving The Web: The Past, Present and Future of the World Wide Web by its Inventor* (Texere 2000).

⁴⁶ R. Ascierto, ‘Senate Mulls Fate of Net Neutrality’ *ComputerWire* (14 June 2006).

⁴⁷ Mark Lemley and Lawrence Lessig, ‘The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era’ 48 *UCLA Law Review* (2001) 925, at p. 956.

⁴⁸ Testimony of Gary R. Bachula before the Senate Committee on Commerce, Science and Transportation Hearing on “Network Neutrality” (7 February 2006) p. 2 available from: <http://commerce.senate.gov/pdf/bachula-020706.pdf>.

6. Ofcom's attitude to net neutrality in the UK

In the UK, Ofcom's principal duties in carrying out its regulatory functions are twofold: It must "further the interests of citizens in relation to communication matters" and "further the interests of consumers in relevant markets" by promoting competition.⁴⁹ In particular, Ofcom is tasked with ensuring "the availability throughout the United Kingdom of a wide range of electronic communications services".⁵⁰ It is also expected to have regard to, amongst other things, "the desirability of encouraging investment and innovation" and "the desirability of encouraging the availability and use of high-speed data transfer services".⁵¹ The various dichotomies encapsulated in this role – consumers/citizens, innovation/investment – perhaps explain why Ofcom's has, to date, adopted a "wait and see" attitude towards net neutrality regulation.

In many ways Ofcom's position is perfectly understandable. One of the key drivers underpinning the debate in the US is the perceived lack of competition in the US market for broadband – cable or DSL services dominate the market.⁵² With limited options for consumers to switch, proponents of net neutrality maintain that there are no countervailing market forces in the US to curb discriminatory actions. In contrast, the retail broadband market (as opposed to the wholesale market) in the UK does not suffer from a lack of competition. Indeed most players in that market would describe it as cut-throat and many are struggling to generate any profits from broadband services at all. As a result the operators' investment arguments described above are perhaps harder to refute in the UK.

Whereas Ofcom has yet to issue any definitive statements on the issue, there are suggestions that it regards countervailing market forces and customers' ability to switch as adequate protections against the need for affirmative net neutrality regulation.⁵³ Hence we may see regulation "around the edges" with a basic consumer protection flavour. These measures may include the provision of information regarding services and applications that will or may be degraded or unavailable

under certain conditions.⁵⁴ Likewise, facilitating customer switching quickly and inexpensively – through regulated contractual termination clauses – may be considered. Indeed in August 2006, Ofcom proposed that rules designed to make it easier for customers to transfer between broadband providers be introduced. These provisions are intended to make it easier for customers to obtain a Migration Authorisation Code (MAC) – required to switch providers – from their current supplier. Regulators may also take a more direct role in the market by promoting (or reducing barriers to entry for) other forms of access such as BPL or publicly funded wireless broadband in areas with high population densities.

7. A way forward?

None of the regulatory measures being considered by Ofcom would in themselves prevent access tiering. Instead they would be designed to make tiering (i) more transparent to customers and (ii) less appealing to operators. Some may question whether this goes far enough. In the net neutrality debate the interests of a large body of Internet users are under-represented. Operators and content and service providers are able to make their voices heard – as they have in the US. But there is a "third voice" to consider, those individuals and entities who have embraced the read/write Internet and whose main interests relate to a broad range of matters in the realm of "free speech" or "individual autonomy" or "democratic participation". Those who use the Internet not just as a communications medium, but as a productive medium as well. These interests have been able to flourish in the open-access environment of the Internet circa 1999, but are likely to be subsumed by the access tiering proposed by operators. This is the voice that is likely to suffer if access tiering becomes wide spread and broadband capacity does not meet demand. This component of the debate strikes a chord with Ofcom's duties towards citizens, quite apart from its duties to consumers. It also presents a set of interests that may be under-represented in some competition law discussions on the issue of net neutrality that tend to focus on pricing issues.

In the UK, the BBC recently announced, as part of its Creative Future editorial blueprint, an increased focus on user generated content. "Audiences of all ages", the BBC states, "not only want the choice of what to watch and listen to when they want, they also expect to take part, debate, create and control. Interactivity and user generated content are increasingly important stimuli for the creative process."⁵⁵ This is the crux of the public policy question. Visualising the Internet in static content distribution-type terms misses that end users

⁴⁹ Communications Act 2003, section 3(1).

⁵⁰ Communications Act 2003, section 3(2)(b).

⁵¹ Communications Act 2003, section 3(4)(d)–(e).

⁵² As of the end of 2004, the FCC reported that incumbent cable carriers and incumbent local telephone companies accounted for approximately 93% of all broadband access to homes and small offices: http://www.fcc.gov/Bureaus/Common_Carrier/Reports/FCC-State_Link/IAD/hspd0705.pdf. Similarly, the US Government Accountability Office found that the median number of broadband providers available to a given household is two. See Government Accountability Office, 'Broadband Deployment is Extensive throughout the United States, but It Is Difficult to Assess the Extent of Deployment Gaps in Rural Areas' (May 2006) p. 23 available from: <http://www.gao.gov/new.items/d06426.pdf>.

⁵³ See Ofcom presentation delivered by Dougal Scott on 25 May 2006 at the London Business School entitled 'Net Neutrality in the UK' (slide 13) available from: http://www.london.edu/assets/documents/PDF/Dougal_Scott.pdf.

⁵⁴ In this context we note that the UK ISP Association is currently drawing up a non-binding "best-practice statement" aimed at encouraging its members to clearly state whether certain types of traffic are subject to access tiering. See 'UK ISPs hope ground-breaking guidelines will leave net neutrality up to market forces' *Telecom Markets* (30 May 2006).

⁵⁵ 'Creative Future – BBC addresses creative challenges of on-demand' (BBC Press Release, 25 April 2006) available from: http://www.bbc.co.uk/pressoffice/pressreleases/stories/2006/04_april/25/creative.shtml.

on the Internet are not just end points in the content distribution chain; rather they are pivotal figures in an increasingly immersive media environment. Logically, it is these end users who should be the focus of any regulatory enforcement of net neutrality.

One option, therefore, is for tiering choices to sit with users rather than operators (or, indeed, regulators). This would help ensure that deviations in network performance are based on the needs of end users rather than the needs of particular application providers. Measures to provide end users with that choice, which would operate in tandem with the disclosure and contract termination requirements mentioned above, might include:

- Restricting operators from discriminating amongst particular application providers, whilst allowing them to provide different broadband packages with varying sets of quality of service guarantees for particular types of application selected by the user.
- Requiring operators to give a clear description of any differences between download and upload capabilities. Not in the sense of “5 Mb/s download, 200 Kb/s upload” but in terms of meaningful information about upload performance for a variety of services – podcasts versus video streams, say. This could be coupled with the requirement to offer minimum upload speeds (possibly at the expense of download bandwidth) at the request of the user and where technically feasible.
- Prohibiting operators from censoring the content, or sources of content, which travel across their networks. Censorship at some level is obviously required to deal with very real problems such as child pornography, spam and viruses, but that is a very different thing to censoring content based on commercial or political interests. We are comfortable with dealing with “free speech” protections (and the limits thereon) in the offline environment and there is no reason why we should not be able to do the same online.

If these measures are seen as too restrictive on operators, the next best alternative may be to adopt something akin to a compulsory licensing regime coupled with minimum access speeds and service guarantees. In this scenario, rates for enhanced quality of service would have to be applied universally

irrespective of the type of data that the content or service provider wishes to transport. Non-commercial participants would be protected by the minima, which might require a source of funding – a “communications tax” – akin to the BBC licence fee, although this would likely face considerable opposition itself and create a bureaucratic Internet framework which would present its own set of problems.

8. Conclusion

The net neutrality debate throws up fundamental questions about the structure and form of our existing and emerging communications and content production environment – the centre of which is the Internet. It is the answers to these questions that should inform and structure any regulatory response (or lack thereof).

Jonathan Zittrain argues that the Internet “is at a point of inflection”.⁵⁶ Security concerns allied with stability and reliability requirements have placed extraordinary pressure on the Internet’s openly accessible end-to-end design. End-to-end as a design philosophy must, therefore, make concessions to practical reality. Zittrain argues that “Strict loyalty to end-to-end neutrality should give way to a new generative principle, a rule that asks that modifications to the PC/Internet grid be made when they do the least harm to its generative possibilities”.⁵⁷

Consumer-led, rather than operator-determined, access tiering, matched with meaningful disclosure requirements and contractual protections best balance the reasonable demand that there be an incentive to invest in Internet infrastructure with the public interest in a “non-discriminatory” Internet. Net neutrality requirements in their strongest incarnation, whilst laudable in many respects, are not practical or desirable. Giving consumers full control of the shift from a “neutral” Internet to a tiered space may be one way of minimising the negative effects that this forced migration will inevitably produce, while enabling the benefits of this technological marvel to be fully realised.

Paul Ganley (paul.ganley@bakernet.com) and **Ben Allgrove** (ben.allgrove@bakernet.com) are both associates in the IT/Com department at Baker & McKenzie LLP in London.

⁵⁶ Jonathan Zittrain, ‘The Generative Internet’ 119 *Harvard Law Review* (2006) 1974, at p. 2039.

⁵⁷ *Ibid.*, at p. 2031. “Generativity” is Zittrain’s description of the technological design of the PC/Internet “grid” that promotes application and content layer innovation by a vast range of Internet users in various technical and creative fields.