CONSULTATION RESPONSE: TRAFFIC MANAGEMENT & NET NEUTRALITY

Submitted by Dean Bubley, Director & Founder, Disruptive Analysis, Suite 13, 19-21 Crawford Street, London W1H 1PJ, UK

1. Introduction

Disruptive Analysis a London-based industry analyst and consulting firm, covering the telecoms industry. In recent years, it has had particular focus on new mobile technologies and the evolution of broadband traffic management, and in June 2010 published a study covering the key techniques being used and evaluated by mobile operators.

It acts an advisor to a range of technology vendors, operators and industry bodies; its founder Dean Bubley regularly comments on related issues on his blog, in discussions with journalists, and at industry conferences.

2. General comments

As an over-riding comment, Disruptive Analysis believes that discussions around Net Neutrality and Traffic Management should distinguish carefully between:

- The Internet and the provision to customers of Internet Access Services
- Other data services transiting broadband access lines

There is a strong argument for *Internet* Neutrality, where the service sold is billed as *Internet* Access. This should not preclude fixed or mobile operators from also offering other non-Internet services over the same physical broadband access connections – for example, their own IPTV, a corporate VPN service for home-workers, or a connection municipal smart-metering service.

Traffic management issues can thus broadly be divided between:

- Discrimination between Internet Access and other services, on the same broadband access connection
- Discrimination *within* the Internet Access Service, between different websites or applications

As a general principle, Disruptive Analysis believes that *within* the Internet Access offering, there should be strong limits on the ability to conduct prioritisation, degradation and other forms of intrusive traffic management.

But there should also be the ability for other services to take priority over general Internet access traffic on the access network, as long as this is made clear to the end user – and, *crucially*, the end-user has a reasonable expected level of quality for that service. It is also considered *unreasonable* to block or degrade Internet traffic (or

parts thereof) in the absence of any congestion, but merely for commercial reasons related to specific Internet applications.

In a way, the term "Internet" ought to have a similar status to a French wine's "Appellation d'origine contrôlée". To extend the analogy – if a telecom operator wishes to sell "Champagne", it must be the real thing – but that does not preclude it from selling Cava as well. In this context, the "Internet" involves as direct-as-possible connection from the user's exchange or radio base station to an Internet exchange or peering point, transiting the operator's backhaul, transport and core network infrastructure with minimal interference. Something going via an operator-controlled server that makes substantive changes (eg compression of video traffic, or filtering of particular applications) is not "pure" Internet access.

It should be possible for broadband operators to work with specific Internet companies to develop non-Internet, customised services, such as an enhanced variant of YouTube or BBC iPlayer, which *transits the operator's own application servers before it then connects to the Internet.* This could be considered an optimised version of Internet content or applications, and a service in its own right, perhaps saleable separately from basic Internet Access. A given user might therefore be able to receive both "native" (pure Internet) YouTube or an optimised/prioritised one.

It is also important for operators and ISPs to start developing a "Congestion API", which gives data on the real-time (or near real-time) condition of the network, to users, application providers and other third parties. This would ensure that unfair restrictions are not placed on traffic when there is ample capacity available – and would also allow all parties to make sensible decisions when congestion does occur (eg. pay for some form of priority, accept poorer-quality applications, wait until another time etc).

This is similar to many examples in normal life – eg digital road congestion signs suggesting alternative routes (or "expected delays" because of predictable events).

Where resources are used for both Internet access and non Internet services (eg home broadband lines, cell-site backhaul) and in particular where they are shared between multiple users, things get more complicated. In theory, it is OK with best-efforts "full Internet access" being at the bottom of the pile, after other services have been dealt with. But there needs to be clarity in marketing, management and oversight. To use an analogy - if you buy a standby ticket for a flight, you know you might not get on. But if the plane takes off with empty seats and you're still refused boarding, you've got grounds for complaint & redress. Best-effort means that the service provider *really* has to make their "best effort" to accommodate you - *and* is still subject to the original contractual terms and subsequent reporting.

Given the benefits of the Internet to society and economy, it may be important to mandate that all operators offer basic unfettered Internet access service, as a pre-condition for being able to offer other access services as well. This is an area for potential future study.

3. Diversity of Traffic Management options

In the past two years, there has been an astonishing proliferation of different technical approaches to broadband traffic management. Ofcom has identified some of these in its discussion document, while Disruptive Analysis conducted a somewhat broader review in mid-2010.

Virtually every vendor in the marketplace has recast their products into a solution intended to help manage the data traffic explosion. Offload to femtocell or WiFi, better policy-management, compression or transcoding of "rich" data and video, many forms of optimisation techniques – all are positioned as the panacea for networks struggling under the demands of laptops and smartphones.

Each of these propositions is aimed at different constituencies within the operators – the radio network planners, those implementing new IP cores, others concerned with marketing or user experience and so on. As each of these silos has its own problems and bottlenecks, it is quite common – and entirely understandable - for decisions to be made on narrow criteria, without reference to the overall end-to-end system.

And yet nowhere is there a collective analysis of *all* the techniques in the same place, let alone any view on what a *holistic* strategy for mobile broadband traffic management might be. The danger of deploying tactical, "local" solutions throughout an operator's domain is that they may attempt to fix the wrong problems – or, worse, lead to damaging unintended consequences elsewhere. This is particularly critical, as many operators appear to be following a process of short-term fixes to fix immediate pain – with a view that more strategic approaches might follow in 18-24 months' time.

Figure: Top 10 technologies for mobile data traffic management

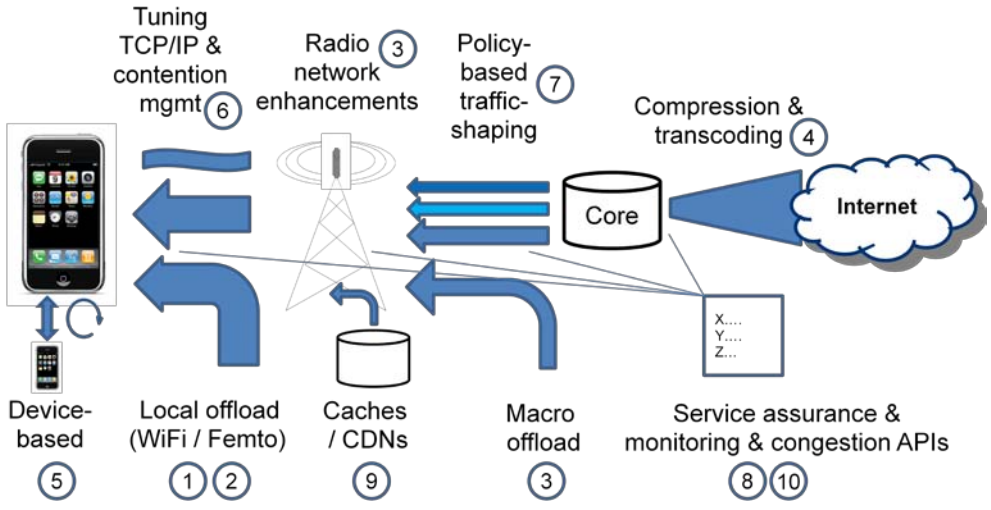

Source: Disruptive Analysis, May 2010

The correct mix of these approaches will vary between different mobile operators – dependent on their existing network architecture, user base, local regulation and marketplace and vendor preferences. There will also be a broad set of individual

challenges around evolving business models and the telco's own internal organisational dynamics.

Further details on each of these approaches is available on request.

4. Specific Questions

Question 1: How enduring do you think congestion problems are likely to be on different networks and for different players?

History suggests that data traffic tends to expand to fill the space available, on both fixed and more recently mobile networks – as long as appropriate pricing and device availability permit.

In the fixed broadband world, there is strong possibility of continued strong growth of data traffic, especially with the advent of Internet-enabled TVs which could stream/download HDTV or 3DTV to living-room scale screens. Potentially, these can consume 5GB of data per hour or more, which may be hard to supply economically – especially as TVs tend to be left switched-on and displaying images, even if there are no viewers present. That said, continued rollout of NGA based on fibre should be able to handle significant traffic volumes, although careful focus will need to be placed on the pricing.

The mobile domain is very different, as there are both physical constraints (amount of spectrum), practical constraints (cell-site availability and planning permission) and unpredictable factors (indoor coverage and user mobility). Capacity is maximised by giving the largest amount of radio resource to those with the best signal conditions – not those users or applications deemed the most important. These make guarantees of non-congested networks extremely difficult, even with the advent of LTE, new frequency bands and innovations such as femtocells.

The timelines for capacity expansion (especially in times of capital constraint) are very challenging given the extremely rapid adoption of smartphones and PC mobile broadband. The addition of new device categories such as tablets (like the Apple iPad) and the probability of sudden and unpredictable "viral" adoption of new applications on open operating-systems will make network planning extremely complicated – congestion is inevitable in some circumstances.

It is also important to note that mobile network congestion can be caused by factors other than just traffic volume (MB or GB of data). In particular, there is evidence to suggest that the excessive radio signalling traffic generated by some devices and chipsets can impact the *effective* capacity, by impacting the control elements of the network, rather than saturating the bandwidth with excess traffic. In future there may be other bottlenecks exposed by ongoing trends in application and device innovation – managing "traffic" is not necessarily synonymous with managing "congestion"

Question 2: What do you think are possible incentives for potentially unfair discrimination?

- Desire to block or degrade the quality of Internet applications seen as competitive to an operator's own services (eg Skype, Facebook)
- Desire to prevent a rival operator from "offloading" mobile data traffic onto their network via WiFi or femtocells
- Desire to make open-Internet applications and services less attractive than the operator's own hosted variants
- Attempts at modifying user behaviour to reduce overall traffic loads

Question 3: Can you provide any evidence of economic and or consumer value generated by traffic management?

Not answered

Question 4: Conversely, do you think that unconstrained traffic management has the potential for (or is already causing) consumer/citizen harm? Please include any relevant evidence.

- Potential for reducing the attractiveness of popular (but data-heavy) services such as YouTube or Facebook
- Potential for decreasing mobile handset battery life, for example by artificially slowing download speeds for videos, thereby keeping the radio in active state for longer.
- Potential to make it hard to assess the quality of new services
- Potential to limit investment by new startups unable to compete with established application providers in working with operators

Question 5: Can you provide any evidence that allowing traffic management has a negative impact on innovation?

At present, there is no evidence of this – although it is noteworthy that those countries that permit blocking of VoIP services have not produced any major software innovators in that field domestically.

Question 6: Ofcom's preliminary view is that there is currently insufficient evidence to justify ex ante regulation to prohibit certain forms of traffic management. Are you aware of evidence that supports or contradicts this view?

The best-known example is that of the "collateral damage" occasioned by Comcast's blocking of P2P traffic in the US, which also apparently affected some business users' Lotus Notes transmissions as an unanticipated side-effect.

The risk of "false positives" such as that must be born in mind when considering any benefits of traffic management, especially where the network attempts to discriminate based on "application profiles" which may be inaccurate or poorly-targeted.

Question 7: Ofcom's preliminary view is that more should be done to increase consumer transparency around traffic management. Do you think doing so would sufficiently address any potential concerns and why?

In general, Disruptive Analysis agrees that sufficient information, coupled with sufficient competition *and ease of switching* should reduce the need for legislation about traffic management. As long as customers *and application providers* have the ability to understand and monitor traffic management, the market should operate efficiently.

It is important that "upstream" companies such as web and media companies are able to discern traffic management policies and congestion status of networks, so that they can inform their users that any examples of poor service can be attributed to their access provider, and not themselves.

One can imagine a web video provider alerting a user that "*there is congestion on your service provider's network and it is likely to reduce the quality of the streams you are receiving – would you (a) like to continue anyway, (b) pay extra for higher quality or (c) consider switching to another service provider such as our partner ABC*"

Question 8: Are you aware of any evidence that sheds light on peoples' ability to understand and act upon information they are given regarding traffic management?

The general shift away from mobile data "walled gardens" towards more generic Internet access could be interpreted to mean that users can understand, and vote with their wallets, when offered the choice between unmanaged and managed services.

Question 9: How can information on traffic management be presented so that it is accessible and meaningful to consumers, both in understanding any restrictions on their existing offering, and in choosing between rival offerings? Can you give examples of useful approaches to informing consumers about complex issues, including from other sectors?

Firstly, there needs to be clear labelling and advertising of services, in particular making a distinction between "the Real Internet" and other Internet-like services which may be filtered, prioritised or otherwise interfered-with.

A good analogy is the food industry, which has various terms like Appelation Controlee, Certified Prime Beef, types of chocolate and various regional brands such as Champagne, Cornish Pasties.

There needs to be a distinction made by operators between "pure" Internet Access, vs. "Processed Internet-type substitute" or "Internet-flavoured access". Clearly, there will need to be a separate debate about what "certified Internet" is – but something based around accepted industry standards and protocols should be a starting point.

Once a clear distinction is made between Internet access and other services, there then needs to be some clarity on how traffic management *between* these categories

works. Given the broad set of parameters (upstream and downstream speed, use of signalling resource, buffering, degradation etc) it is unlikely that full details can be provided clearly in a retail environment, although there is no reason to hide the full details online.

Another possible comparison or analogy might be to compare subscribing to broadband Internet access (fixed or mobile) as similar to renting a property. The landlord should allow you "quiet enjoyment" without undue interference or limits, but is permitted to enter in emergencies, specify upfront rules (no pets) and hold you to account for destructive behaviour (withholding a deposit)

Question 10: How can compliance with transparency obligations best be verified?

Various organisations already produce monitoring tools, which can be used to track the performance of network connections and particular applications / traffic types. Given a sufficient sample size and data analysis capabilities, it should be possible for non-operator organisations (commercial or regulatory) to assess the prevalence, nature and impact of traffic management.

For example, a body such as Google could probably "reverse engineer" traffic management policies by examining the network performance of millions of Android smartphones users. Any "unnatural" patterns not explainable by normal variation of network conditions should be relatively easy to discern – eg degradation at particular times of day, or for specific websites within a group of peers.

Such 3rd-party monitoring tools should be encouraged as a means of ensuring that competition and transparency are maintained. However, it may become necessary to protect such test/measurement data traffic from interference itself.

Question 11: Under what circumstances do you think the imposition of a minimum quality of service would be appropriate and why?

At the moment, many mobile broadband services are sold on the basis of unachievable peak-rate speeds for data download. There is a strong argument that operators should be made to quote "realistic expected averages" under normal use (however defined) and then made to publish actual figures to show how close they reached those expectations.

It is also very important to consider not just "quality of service", but a broader view of "quality of experience" – for example including coverage as well as speed.