



VIEWS AND NEWS

ECONOMICS AND TECHNOLOGY, INC.

September 2010

IN THIS ISSUE

- Canadian regulators expand requirement for nondiscriminatory provision of wholesale broadband transmission to rival carriers
- Net Neutrality, “Paid Prioritization,” and “Network Management” – Part I

Canadian regulators expand requirement for nondiscriminatory provision of wholesale broadband transmission to rival carriers

Canadian regulators (the Canadian Radio-television and Telecommunications Commission – CRTC) have consistently treated high-speed Internet access as a common carrier telecommunications service, whether provided over ILEC facilities or over those owned by cable companies. In 2006 and 2007, the CRTC adopted a requirement that Canada’s larger ILECs and cable carriers make high-speed Internet access services available to competitors at whatever speeds the incumbents offer to their own retail Internet customers – the so-called “speed-matching” rule. After ILECs and cable carriers mounted challenges to the speed-matching requirements – particularly as they applied to higher-speed services – the CRTC initiated a proceeding to examine the current status of wholesale and retail Internet access competition to determine what changes, if any, were appropriate. ETI had submitted expert evidence in that proceeding on behalf of the largest CLEC in Canada. The CRTC has just released an order that not only reaffirmed the earlier ruling, it even broadened the speed-matching requirements with respect to incumbents’ high-speed Internet access services. The relevant findings upon which the CRTC based its August 30, 2010 decision include:

- The additional competition supplied by non-facilities-based retail providers of Internet access services is beneficial to residential and small business customers:

“The retail residential and small-to-medium-sized business Internet service markets are now served by the incumbents and a number of smaller competitors that generally use the incumbents’ wholesale services to do so. In the Commission’s view, these competitors’ services bring pricing discipline, innovation, and consumer choice to these retail Internet service markets.”

- Retail competitors are impaired without access to wholesale ILEC and cable carrier high-speed Internet access services:

“[W]ithout a speed-matching requirement for wireline aggregated ADSL access and TPIA [Third-Party Internet Access] services [the high-speed wholesale Internet access offerings of ILECs and cable carriers, respectively], it is likely that competition in retail Internet service markets would be unduly impaired. In the Commission’s view, an ILEC and cable carrier duopoly would likely occur in the retail residential Internet service market, and competition might be reduced substantially in small-to-medium-

sized retail business Internet service markets.”

“[A]t this time, retail Internet services provisioned using wireless and satellite facilities generally remain complements to, and not substitutes for, retail Internet services provisioned using wireline facilities.”

- Calling a service or facility “next generation” (or “broadband”) does not confer some mystical status that trumps the regulators’ obligation to examine the impact of denying wholesale access on the viability of retail competition:

“[T]he real issue is to establish those wholesale obligations, if any, that should apply to identified facilities. ... competitors continue to require access to the wholesale services currently offered by the incumbents over their digital subscriber line and DOCSIS platforms in order to ensure that sufficient competition exists in the provision of retail Internet services. In the case of the ILECs, the facilities that are subject to wholesale obligations include FTTN and, in the case of the cable carriers, DOCSIS 3.0 facilities.”

In stark contrast to the Canadian approach, for most of the past decade the FCC has classified broadband Internet access as an “information service,” thereby relieving US ILECs and cable providers of any obligation to make the underlying transmission component available to competing retail ISPs on a wholesale basis. Like the FCC, the CRTC does not impose rate regulation on retail Internet services, a policy that dates back to the time when “competitive retail Internet service providers were able to provide their services on a dial-up basis, using the customer’s retail telephone service” and the ILECs and cable companies were not themselves heavily involved in providing retail Internet services. However, unlike the FCC, the CRTC adapted its policies to the changes in the Internet access market once the dominant facilities-based telephone and cable providers began to offer broadband Internet access. Thus, “[a]s the retail Internet service market began to evolve to higher speed retail Internet services, to ensure these services remained subject to competition sufficient to protect consumer interests, the [CRTC] required that ILECs and cable carriers make some of their high-speed access facilities available as wholesale services for competitors to use as inputs in the provision of retail Internet services.” *CRTC 2010-632* at para. 6. The CRTC has been willing to adopt and maintain regulatory forbearance for retail Internet access services precisely because it has a wholesale regulatory framework that assures non-facilities-based ISPs the ability to purchase wholesale services –

including broadband – from ILECs and cable providers.

As a result of the affirmative measures that Canadian regulators have taken to promote retail competition by imposing wholesale obligations on ILECs and cable companies, Canadian regulators have reduced the scope of net neutrality concerns that remain to be policed through *ex post* (after the fact) regulation (such as is contemplated by the FCC's "Third Way" proposal – see *Views and News*, June 2010). Having adopted requirements that foster competition at the retail level, the CRTC's net neutrality framework (see, Telecom Policy CRTC 2009-657, *Review of the Internet traffic management practices of Internet service providers*, 21 October 2009) legitimately presumes that retail ISPs are implementing their Internet traffic management practices (ITMPs) in a competitively neutral manner (although such practices remain subject to *ex post* review in the event of a customer complaint, and "transparency," i.e., full disclosure of ITMPs, is still required). By contrast, ITMPs that a "primary" (ILEC or cable carrier) ISP seeks to apply to its wholesale services are subject to advance regulatory scrutiny. For wholesale services, any economic ITMP must be disclosed in a filed tariff, which is subject to prior approval by the CRTC "using the ordinary principles for rate approvals." Some technical ITMPs can be implemented without prior approval of the Commission, if they are not more restrictive on "secondary" ISPs (facilities-dependent providers) than on the primary ISP's own retail Internet services and "do not have a significant and disproportionate impact on secondary ISP traffic." But when a technical ITMP is more restrictive to retail competitors than to the primary ISP's own retail service, the CRTC must review and approve the provision prior to its implementation.

The CRTC has thus crafted a regulatory model that, by imposing greater regulatory controls at the wholesale level, helps to facilitate increased competition; in so doing it has largely eliminated the need for regulatory involvement at the retail level. If a retail ISP were to unilaterally undertake to impose an ITMP whose effect is to disadvantage its retail customers (by, for example, degrading or blocking their access to websites or services that are not willing to play by the retail ISP's ITMP rules), the affected customers can "vote with their feet" and seek out a competing provider whose policies are more open and neutral. While not necessarily a perfect solution, the CRTC regulatory model imposes regulation in those upstream segments where market failure could potentially arise precisely so as to reduce the potential for such market failure in downstream segments that rely upon the upstream inputs. In contrast, the FCC's decision not to impose common carrier obligations on upstream facilities-based services has created conditions for downstream retail services that make it nearly impossible for any independent non-facilities-based retail competitor to enter and remain viable over time.

For more information on this subject, please contact Helen E. Golding at hgolding@econtech.com.

Net Neutrality, "Paid Prioritization," and "Network Management" – Part I

Among the most contentious issues in the "Net Neutrality" debate are those that revolve around the right of Internet access providers to charge customers for prioritizing certain traffic and otherwise to engage in what they describe as "network management"

practices whose effect would be to limit or prevent outright certain uses of their Internet access services. As the debate has intensified in recent months, the various positions on these subjects generally fall between these two bookends:

- Certain public interest organizations and Internet purists want service providers to be prohibited from examining or affording any special treatment to Internet traffic based upon the *content* of the transmission. They oppose any form of "paid prioritization" and content-based "network management." To them, such practices are an anathema to a truly "open" Internet, and they will be satisfied with nothing short of an outright prohibition of any such practices.
- The service providers' position is that it is *their* network and that they are entitled to run it as they see fit. This view is perhaps best summarized in a remark made several years ago by former SBC/AT&T CEO Ed Whitacre, when he said: "How do you think they're going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent the capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using."

This is the first of two articles in which we explore those aspects of each side's argument, perhaps as a basis for ultimately reaching common ground. This month, we examine "paid prioritization" and the closely-related "managed services" issue. In October, we will address the controversy regarding "network management."

"Paid prioritization"

One of the most controversial aspects of the net neutrality debate centers around whether Internet service and backbone network providers should have the right to offer customers prioritized delivery of Internet traffic and to impose additional charges for such prioritization services. The proposed "prioritized delivery" would run all the way to the end user "eyeballs," with the end user's ISP offering to prioritize a content/application providers' traffic to/from a given user for a fee to be paid by that host service provider.

Package delivery services routinely offer several "classes" of delivery service – e.g., overnight, two-day, three-day, 7-day, etc. – charging the shipper more for expedited services. By offering an array of delivery choices at different prices, customers are able to make economic trade-offs between the value and importance to them of prioritized delivery vs. the additional costs that such prioritization entails. Shouldn't providers of Internet transport and delivery services be able to offer similar delivery options?

Indeed, setting aside, for the moment, the value or importance of the specific information being transmitted, prioritization affects different types of data packets differently. For example, packets supporting a real-time interactive exchange of information, such as a Voice over Internet Protocol (VoIP) telephone call, are far more sensitive to the effects of circuitous routing of packets and so-called "latency" – where packets arrive at their intended destination out of order – than are packets associated with the transmission of e-mails or other text or static images, where

millisecond-level delays required for reassembly of out-of-order packets or for retransmission of incomplete or missing packets are invisible to the end users. Two-way video conferences may present even less tolerance of such transmission issues than voice-only applications. And within each generic category of packets, some may, like their physical package counterparts, present greater value and importance to the sender and/or recipient than others.

For example, in a September 15 *ex parte* submission in the FCC's *Open Internet NPRM*, AT&T posited that

... to improve patient welfare and reduce unnecessary in-hospital expenses, a hospital may, in concept, wish to provide a medical application (e.g., wireless heart monitoring) delivered to patients in their homes over their Internet access connections. Such a monitoring application may have stringent end-to-end network performance requirements that necessitate traffic prioritization, including in the "last mile" connection to the end user. The hospital may wish to purchase such traffic prioritization capabilities from a patient's ISP and transparently include them in the monitoring application, rather than requiring the patient with the heart condition to take responsibility for the technically sophisticated process of acquiring and configuring the requisite prioritization capabilities.

It seems, however, that while intensely controversial, the concept of "paid prioritization" may not be all that well understood. Internet users routinely make choices as to the bandwidth capacity of the service they purchase, with the price increasing at higher bandwidths. End users can select bandwidths ranging from 56 kbps dial-up to broadband offerings ranging up to 50 mbps in one or both directions. Satisfaction of the needs of the patient described in AT&T's *ex parte* letter is likely far more sensitive to the bandwidth selected *by that patient* for the last-mile connection to her home than any packet prioritization that the hypothetical hospital might purchase from the patient's ISP to move the hospital's packets to the front of the line. The range of bandwidths available to application/content provider host websites is far greater than that offered to residential subscribers. Sites that deal with massive volumes of traffic – like amazon, facebook, ebay, twitter, google, etc. – obviously could not meet their users' traffic demands with less than adequate bandwidth – and certainly these high-traffic sites will require immensely more bandwidth than, say, a neighborhood pizzeria or a local public library.

Indeed, many who oppose the concept of paid prioritization hold that the *only* issue is bandwidth. As they see it, each party – eyeball or host – is responsible for selecting and paying for the bandwidth sufficient to meet their needs. Under this view, any requirement to prioritize packets arises only where overall bandwidth capacity is not sufficient to afford all packets immediate and undelayed transmission; thus, where any sort of blockage is present, the correct solution is to expand capacity. Proponents of this view hold that the existence of insufficient capacity is simply a contrived condition that permits the provider to degrade certain users' service in order to "sell" a higher priority service to other users. Put differently, when someone is allowed to purchase a spot "at the front of the line," that simply pushes everyone else further back.

There is no question that providing sufficient bandwidth to meet an end user's needs offers the most direct means of resolving prioritization concerns. However, one problem with this "no-delay" capacity theory is that, at least in principle, it could lead to a highly

inefficient network, where costs could escalate without producing very much in the way of incremental benefit. Suppose, for example, that a package delivery company were to adopt this same sort of "no package left behind" business model, and as such provided guaranteed overnight delivery to every package without any specific overnight delivery charge. To accomplish this, the carrier would need to add more airplanes to handle the increased volume of packages requiring overnight transport and, while no longer imposing an overnight delivery service charge, the average cost of *all* packages handled by the carrier – including those formerly 2-, 3- and 7-day deliveries that would otherwise have gone by truck or train – will necessarily increase. Yet for many customers, the 2-, 3- or 7-day schedule would have been quite adequate, with the overnight service affording them little or no incremental benefit or value. Offering different priorities at different prices creates a direct linkage between the cost of the delivery and the value of the delivery to the shipper, and allows each shipper to make an intelligent economic trade-off between cost and speed.

Of course, this presupposes that there is some reasonably close relationship between the *price* that is being charged for each class of delivery service and the *cost* that the carrier confronts in providing each class of service. Since the market for package delivery services is competitive, the structure of prices will come to reflect the structure of costs, because if any one of the competing carriers seeks to extract an excessive above-cost price for its high-value overnight delivery service, its rivals can be expected to undercut that price and in so doing capture additional business.

And therein lies the problem when it comes to broadband Internet access. Here, the market for connection to end user "eyeballs" is not even close to being competitive. As the FCC has recently noted on several occasions, consumer broadband Internet access is provided by a duopoly consisting of the local telephone and cable companies, and that price competition is not likely to develop in this type of market structure. Indeed, earlier this year the FCC issued its *National Broadband Plan*, which noted that "[g]iven that approximately 96% of the population has at most two wireline providers, there are reasons to be concerned about wireline broadband competition in the United States." Each of the two service providers has considerable market power. As such, each can be expected, if permitted to do so, to utilize "paid prioritization" as a device for segmenting the market on the basis of value of the content being transmitted – extracting monopoly rents from hosts and/or eyeballs based on the perceived value of the content, without regard to the actual cost of the Internet access service being furnished. This is a legitimate concern, one that may transcend the potential inefficiencies that might arise without paid prioritization.

Absent the pressures of a competitive market, the non-competing duopolists would be free to segment the market in a manner that is far more likely to be driven by the value of the content of the transmission than by whatever "prioritization" attributes may be provided. Absent an effective competitive threat, there would be nothing to prevent AT&T, in the example provided in its *ex parte* submission, to impose a value-based fee for whatever "prioritization" might be required by a heart monitoring service where, if AT&T is to be believed, even a

millisecond delay could result in a loss of life. More generally, there would be nothing to prevent an ISP from imposing a value-based charge for delivery of video content to its “eyeballs,” charging the content provider more for the transmission of high-value content, such as a major sports or theatrical event than for the transmission of reruns of *Friends*. One must be suspicious of the types of contrived examples that AT&T has put forth as requiring prioritization, such as AT&T’s claim that (most likely unnecessary) “stringent end-to-end network performance requirements,” “traffic prioritization,” and last-mile facilities upgrades are necessary for telecommuting and distance learning, which in turn will “reduce greenhouse gasses.”

That said, at a certain level the availability of paid prioritization services – like their physical analog in the form of expedited package delivery – may in theory contribute to the overall efficiency of the nation’s broadband infrastructure, but only if three specific conditions are satisfied:

- 1) The end user, not the application/content provider, must be the “customer” who specifies and pays for the desired prioritization of his traffic.
- 2) The prioritization offering must be entirely generic, bearing no relationship to the content of the transmission.
- 3) Any charges for such prioritization must be cost-based, either by virtue of having sufficient competition in the “last mile” to assure a competitive outcome or, alternatively, by means of economic regulation of any such fees.

Whatever economic or technical virtues paid prioritization may offer – if any – pale in comparison to the opportunities for exploitation of the last mile providers’ market power, one result of which could be vertical foreclosure in downstream application and content markets. Before proceeding further down this slippery slope, regulators and policymakers need more than the type of vague arm-waving claims as to the need for and benefit of any form of paid prioritization than those “examples” being put forward by AT&T. And only when, as and if the proponents of paid prioritization were able to offer such evidence can a rational, fact-based decision on this issue be crafted.

Managed Services

Closely related to “paid prioritization” is the concept generally referred to as “managed services.” Managed services are furnished by the last mile broadband service provider using dedicated bandwidth earmarked for a specific service, application or content. For example, when a cable company offers “digital voice” telephone service utilizing VoIP technology, it typically sets aside a certain amount of dedicated IP bandwidth – in effect, a separate IP transmission path – physically provided over the same end user cable facility that is used to provide broadband Internet access and video services. By using this type of service configuration, the voice telephone service is segregated both from the customer’s other IP traffic as well as from that of other customers who are also served by the same physical coax segment.

Similar dedicated IP channels may also be established by the service provider for other purposes – a particularly good example is IPTV, but one could envision other applications for which this type of service arrangement might be used, including, for example, the medical monitoring application alluded to in the AT&T *ex parte*.

So what’s the problem with such “managed services” arrange-

ments? There may be no problem at all, *as long as the broadband service provider makes technically comparable dedicated channels available to competing downstream providers of the same types of services that the service provider itself offers utilizing these types of arrangements.* Consider the following example: VoIP type telephone services that are offered using a dedicated channel by the broadband service provider compete directly with what the FCC refers to as “over-the-top” or “nomadic” VoIP services such as Vonage. But when a broadband service subscriber subscribes to one of these competing IP telephone services, the voice IP traffic is co-mingled with the subscriber’s – and (in the case of coax based services) his neighbors’ IP traffic. Where such traffic is sufficiently heavy as to create congestion either within the subscriber’s own Internet access service or in the shared distribution facility, latency and other transmission issues may negatively affect the *quality* of the voice signal. When a dedicated “managed service” channel is used for this purpose, the voice service is isolated from all other IP traffic and is thus insulated from any such congestion-related problems.

This disparity may be of limited (but still some) importance for services whose need for bandwidth is small, such as IP voice telephony. But for applications involving larger amounts of bandwidth, the nonavailability of a dedicated channel may sufficiently degrade service quality of the application or content so as to effectively block downstream entry. Customers are less likely to be attracted to competing IPTV services that are subject to periodic pauses or other interruptions when the broadband provider’s own IPTV offering is not subject to such problems.

Managed service dedicated channels clearly serve a purpose and service providers should by no means be prevented from creating such arrangements. However, any type of dedicated service-specific channel must be defined strictly in terms of its generic technical properties and be offered to downstream application and content providers on a competitively neutral and nondiscriminatory basis. Absent this requirement, downstream competitors will be unable to offer services whose technical quality can be directly compared with the facilities-based provider’s own integrated offerings. In the end, regulation will be needed to assure this outcome, since as long as the last mile consumer broadband services market remains the domain of two dominant incumbents, there will be no competitive pressure that would compel the duopolists to voluntarily make such dedicated channels available to rival downstream providers.

For more information on this subject, please contact Dr. Lee L. Selwyn at lselwyn@econtech.com.

© 2010 Economics and Technology, Inc. All rights reserved.

About ETI. Founded in 1972, Economics and Technology, Inc. is a leading research and consulting firm specializing in telecommunications regulation and policy, litigation support, taxation, service procurement, and negotiation. ETI serves a wide range of telecom industry stakeholders in the US and abroad, including telecommunications carriers, attorneys and their clients, consumer advocates, state and local governments, regulatory agencies, and large corporate, institutional and government purchasers of telecom services.