EX PARTE

Ms. Marlene H. Dortch
Secretary
Federal Communications Commission
445 12th ST, SW
Washington, DC 20554

Re: In the Matter of Broadband Industry Practices, WC Docket No. 07-52;

Dear Ms. Dortch,

The Commission has concluded its third hearing on this docket, and according to press reports, is considering an order. This is to correct some of the factual errors that have infected the record in this matter, to summarize public concern, and to suggest a resolution.

Factual errors concerning TCP Reset

The technique that triggered the controversy over Comcast's management of its Internet access network is the injection of TCP Reset packets into an open virtual circuit. TCP is a connection-oriented, packet stream protocol used by Internet applications that depend on reliable delivery and are insensitive to packet jitter. Each use of this technique closes a single virtual circuit. The debate over its use seems to depend on whether closing particular virtual circuits constitutes a "blockage" or a "moderation" of application traffic. This is an important distinction in light of the Internet Policy Statement's guarantee of consumers' freedoms *to access the lawful Internet content of their choice* and to *run applications and use services of their choice*. The question of timeliness of access is not mentioned in the Policy Statement, so I assume it's meant to be covered by the network management exception and the over-arching requirement of "neutral" management of IP-enabled services access.

If an application is slowed to a rate consistent with reasonable network management, but still allowed access to content and services of the consumer's choice at this rate, the Policy Statement (*arguendo,* taken as a rule) has not been violated. If the consumer's rate of access is nil or below the consumer's desired rate and the rate dictated by reasonable management, there is a violation. The nexus of management techniques and applications is therefore crucial.

At least two witnesses made misleading statements about the effect of virtual circuit disruption on Peer-to-peer applications. There is no question that the use of the TCP Reset has the effect of "blocking" an individual virtual circuit, but witness Jon Peha asserted that blocking individual connections is equivalent to blocking the application that seeks to use them. This can only be true if we assume a one-to-one correspondence between applications and virtual circuits.

In fact, BitTorrent and other peer-to-peer applications open thousands of virtual circuits per application. This can be observed by network sniffing tools such as Wireshark, and in the running code of the Network Address Translators (NATs) that are commonplace fixtures in home networks. Some witnesses stressed the fact that much smaller numbers of these connections carry data at any given time (on the order of four to ten,) but the fact remains that thousands may be open, and active data transfer shifts among them. Peer-to-peer applications are therefore not prevented from running or accessing content by the closure of specific virtual circuits.

Furthermore, peer-to-peer applications are designed to retry and recover from the closure of virtual circuits; they retain limited memory of past successes and failures, and will retry specific virtual circuits regardless of past experience when limited sources of content are available.

Virtual circuits have impacts on network performance. NATs need to keep track of open virtual circuits to provide translation from a home network's external IP address and internal IP address of the circuit's endpoint, and they do this with a software construct known as a “mapping table.” In traditional NATs, this mapping table has a fixed maximum size, and when an attempt is made to open a circuit that would exceed the table's design, the NAT rejects the request. It accomplishes this unpleasant tasks by sending a TCP Reset packet. Many researchers attempting to detect network management by the Sandvine system used by Comcast have actually encountered TCP Resets generated by NATs, and some have admitted the error and apologized[1].

The Internet, through a quirk of history[2], invests network-wide congestion control in TCP, with a mechanism that controls the rate of packet transmission on individual virtual circuits. This mechanism is defeated by applications that use enormous numbers of virtual circuits, such as peer-to-peer applications, and by other applications that simply consume bandwidth more intensely than the norm:

> *To BT's Bob Briscoe, talk of ISPs' unfair congestion-management techniques is misleading, because congestion management on the Internet was never fair. Telling computers to halve their data rates in the face of congestion, as the TCP protocol does, is fair only if all those computers are contributing equally to the congestion. But in today's Internet, some applications gobble up bandwidth more aggressively than others. If my application is using four times as much bandwidth as yours, and we both halve our transmission rates, I'm still using twice as much bandwidth as you were initially. Moreover, if my gluttony is what caused the congestion in the first place, you're being penalized for my greed. "Ideally, we would want to allow everyone the freedom to use exactly what they wanted," Briscoe says. "The problem is that congestion represents the limit on other people's freedom that my freedom causes."*[3]

Reducing the number of open virtual circuits per application and per user account to the range of normal applications imposes fairness on an otherwise unfair system, and make a vital connection that seems to have eluded the first generation of Internet protocol architects: the Internet is not simply a technical system that moves packets from one endpoint to another, it's the underlay of a social system whose ultimate end points are not computers, but people.

When one person generates more network traffic than another, and the amount of that traffic is such as to prevent the free flow of information, the Four Freedoms are in jeopardy unless the ISP takes steps to protect them. My freedom to communicate and to access content and services is dependent not only on the behavior of my ISP, it's dependent on the behavior of every other person using the Internet, and especially those with whom I share the most critical facilities such as neighborhood cables,

---

1 *Rush to judgment: Comcast not blocking web traffic after all*, Eric Bangeman, Ars Technica, April 08, 2008, http://arstechnica.com/news.ars/post/20080408-rush-to-judgment-comcast-not-blocking-web-traffic-after-all.html (retrieved July 25, 2008)

2 The original scheme for congestion moderation, devised in the 1970s, didn't work. It depended on the ability of each router to send a *Source Quench* control message back to the sender of each packet that could not be forwarded due to congestion. Congestion prevented the reliable delivery of these messages, so Van Jacobson devised an emergency method for moderating TCP transmission rate when packets were lost. Further refinements mark packets that encounter congestion without discarding them, such as Explicit Congestion Notification, Internet RFC 3168 (http://www.ietf.org/rfc/rfc3168.txt,) and Re-ECN: http://www.cs.ucl.ac.uk/staff/bbriscoe/projects/refb/.

3 *Internet Gridlock*, Larry Hardesty, Technology Review, July/August 2008, http://www.technologyreview.com/Infotech/20919/page3/ (retrieved July 25, 2008.)

aggregators, and routers.

Consequently, the mere fact that TCP Reset packets are injected into specific virtual circuit does not constitute evidence that the Policy Statement (taken as a rule) has been violated.

Public Concern

To the extent that there is public concern over Broadband Industry Practices, its focus is on freedom of speech rather than network engineering mechanisms. Free Press and its partners have endeavored to show that ISP practices limit free speech, or have the potential to limit it in the future, according to devious plans only visible to skilled practitioners of the Policy Arts. This analysis is simplistic, as I've endeavored to show above.

Free speech depends on the ability of speakers to engage in orderly discourse without undue disruption from any quarter. The Commission has had occasion to observe numerous instances of audience outbursts when certain views were expressed during its en banc hearings. Economist George Ford was jeered and heckled by a Stanford audience emotionally aroused by Professor Lessig's decidedly anti-intellectual graphic images, and I heard jeering when I suggested, at Harvard, that Free Press proposed a *Deadwood System* of network management.

Bursts of high traffic are the technical equivalent of this sort of audience reaction. And just as it's reasonable to silence or eject those who disrupt public meetings, it's reasonable to throttle sources of high network load.

It can also be taken as given that any common network management technique can be represented to a gullible public as an authoritarian practice. Networks are governed by control theory, but human societies had better not be. Networks facilitate social interaction, but only by the movement of information, not of people. Networks are to be managed as natural resources, not as people: we consume network bandwidth in greater or lesser degrees, but we don't consume our neighbors.

Hence, the first step toward reasonable network policy is the abandonment of the anthropomorphic fallacy that packets are participants in a *one-packet-one-vote* system of self-governance. Applying quotas in human systems is always controversial, but it should never be in control-theoretic systems.

We ensure the free flow of information in networks by delivering packets promptly and by managing congestion wisely, not by indulging in anthropomorphic fantasies equating people with mere packets of information.

Resolution of the Conflict

“Neutrality” is a difficult notion to define for share resource networks. Its origin is the realm of traditional public switched telephone networks, where resources are sharply partitioned in such a way that no user creates side-effects on any other user. Hence, its application to shared-resource networks in general and packet-switched networks in particular is especially problematic. The force-fit of this regulatory construct over the Internet is apparent when we consider the important freedoms that are omitted from the Policy Statement.

The Policy Statement describes the user's relationship with the carrier, but not the user's relationship with other users who share network resources. A mature Policy Statement would enshrine the right of network users to be protected from undue restriction of network access by other users.

This is important when we consider the statistical assumptions that guide network provisioning. Commercial Internet access services frequently provide a Committed Information Rate (CIR), a guarantee of minimum bandwidth available at all times. This is provided inside the network by dedicating network resources to the customer in much the same way that the PSTN reserves bandwidth for each caller. The reservation of this bandwidth is accomplished by the purchase of equipment that enables each customer to use the CIR at the same time.

Residential Internet access accounts are much cheaper than CIR accounts because they have no guaranteed minimum bandwidth. As long as residential broadband Internet services have been marketed, carriers have been able to assume a fairly orderly demand for asymmetrical bandwidth, since most Internet usage moves more information from the Internet to the user than in the reverse direction, and most Web access takes place in short bursts.

Peer-to-peer violates every assumption that carriers have been able to make for the last fifteen years. It sends and receives persistently, often for hours or days at a time without break. It pushes more data upstream than it brings downstream, and it has an insatiable appetite for bandwidth. None of these properties is unmanageable, and none need be destructive to the stability or the efficiency of Internet access networks. But they are radically different from the properties of Web access, and there's the rub. Internet access networks can be made to manage peer-to-peer efficiently, or they can be made to manage Web browsing efficiently. Some day we will know how to handle both simultaneously, but we aren't there yet.

The Internet Engineering Task Force has organized a study group on Peer-to-Peer Infrastructure enhancements, but it's just getting off the group and doesn't yet have full status. IETF working groups typically take years to write a specification, and the industry takes months or years to align its practice with it. The Draft Internet Standard on Explicit Congestion Notification mentioned previously, RFC 3168, has adopted in 2001 and is still not a standard feature in Microsoft Windows. So we're probably years away from a fully-specified Draft Internet Standard on the proper management of peer-to-peer protocols on Internet access networks. Consequently, any demand of the Commission to sanction peer-to-peer network management practices for violating Internet Standards is absurd.

Media reports suggest that the Commission may adopt an order which simultaneously forbids Comcast from using its present system of peer-to-peer management and demands information about the banned system. This suggests a failure of the rational process. If the Commission needs data from Comcast in order to evaluate the management system, it can't very well order the system stopped prior to examining the data. And if it has already determined on a basis that can't be fairly described as fact-based that the system needs to be abandoned, why ask for the data? This is tantamount to conducting a hanging and only then a trial.

The demand for data suggests that more study is needed, and I have no disagreement with that assessment. The Commission's uncertainty also suggests that the Policy Statement itself is defective, which is also correct. A mature policy statement would endeavor to provide not just freedoms from carrier misconduct, it would specify responsibilities, such as:

1. The responsibility to keep computers connected to the public Internet free of viruses, bots, and malware that have an adverse impact on the network.
2. The responsibility to identify the priority of outgoing traffic so that the carrier can manage it correctly, or a permission for the carrier to identify it on the user's behalf.
3. The responsibility to use network resources efficiently.

And so on. A list of this sort is less exciting than a list of freedoms, but it's a necessary part of the maturation of the Internet as the common, converged platform for digital communications.

I would also suggest that the Commission need to devote some time the task of developing carrier guidelines regarding reasonable network management, as I outlined in a previous filing. The present system, in which every carrier makes a crap shoot every time he buys a piece of management gear, gambling that the Commission will not look it up on a secret list of approved practices and possibly sanction it, is simply a prescription for litigation. The Commission frequently loses when its punitive actions are reviewed by the courts.

A Reasonable System of Traffic Management

When dealing with link congestion on the Internet access network's first mile, I would advise carriers to adopt a scheme that manages traffic in two stages, to wit:

1. Allocate bandwidth equally across all active accounts of the same type or tier;
2. Prioritize traffic within each user account according to the traffic's requirements.

The first stage protects light and moderate users from denial of access by heavy users, and the second stage protects each user's high-priority applications – such as Voice over IP – from being denied access by the users lower-priority, bulk data applications, such as peer-to-peer. If the user fails to identify traffic requirements, it's appropriate for the carrier to perform this classification on his or her behalf.

It appears that Comcast has implemented a system that performs Stage 2 of this scheme ahead of equipment to perform Stage 1, and has simply tuned it in such a way has to limit the use of applications known to consume disproportionate share of upstream bandwidth, peer-to-peer in seeding mode. This system can produce erroneous results, as petitioner's employees have noted, for users whose seeding is very light.

An order for the wholesale removal of Stage 2 systems would be unhelpful. These systems are an important part of a total traffic management solution, even though they aren't sufficient by themselves. So any order regarding traffic management practices should permit the ordinary management of priority traffic, done in the appropriate way by a component of a more comprehensive system.

One witness is on record complaining about both parts of this system, referring to the “protocol agnostic” Stage 1 as “user history discrimination.” This is an uninformed criticism. The use of systems that allocate bandwidth fairly among all active users, Fair Queuing systems, is a practice of long standing on IP networks, described in Internet Standards in RFC 896[4], in January 1984. There is little doubt of its legitimacy.

Conclusion

I applaud the Commission for holding public hearings on this matter, and for extending an invitation to me to provide testimony at the first. Internet management is a difficult and highly-politicized topic, and a certain element of drama is inevitable around such an issue. The matter cries out for deeper study by the Commission's technical experts and expert consultants from industry and academia. Punitive action is premature, but clearer standards for consumer disclosure, permitted management practices, and user

4 RFC 896 - Congestion control in IP/TCP internetworks, John Nagle, January, 1984. http://www.faqs.org/rfcs/rfc896.html (retrieved July 24, 2008)

behavior are needed.

Yours truly,

/s/

Richard Bennett