

PO Box 3504
Eugene, Oregon 97403

July 15th, 2014

FCC Secretary Marlene H. Dortch
Federal Communications Commission
445 12th Street, SW Room TW-B204
Washington, DC 20554.

Subject: Comments on Proceeding No. 14-28, "Protecting and Promoting the Open Internet."

Dear Secretary Dortch:

Thank you for the opportunity to comment on the proposed "Network Neutrality" rules. The remarks I offer today are provided in my private capacity, and not as a representative of any other organization or group. Unfortunately, I have no option but to **oppose the proposed rules** for the reasons described below.

I. Jitter-, Latency-, and Loss-Sensitive Traffic Must Be Able To Continue To Be Prioritized (And Thus Protected) From Competing Bulk Flows

Voice over IP and interactive video conferencing traffic are well-known examples of jitter-, latency-, and loss-sensitive application traffic. Even seemingly minor delays (or near-negligible levels of packet loss) can significantly degrade the perceived quality of VoIP or interactive video traffic.

However, if VoIP or interactive video traffic are protected via expedited forwarding rules (sometimes referred to as "class of service" or "quality of service" measures), that traffic can easily be sheltered from interference due to other traffic (including aggressive peer-to-peer file sharing traffic or readily-buffered bulk Internet television traffic flows).

Similarly, in-band network control plane traffic, such as BGP routing announcements, should also be able to be given preferential forwarding status. If that's not done, application traffic can interfere with in-band transmission of the critical routing data need to keep the network itself up and functional. Failure to protect in-band network control plane traffic means that the networks' ability to remain functional may actually be put in jeopardy.

Or consider "ping" (ICMP echo) traffic. Most network engineers recognize the value of "ping" traffic as a diagnostic tool, but don't want that tool to be abused. Therefore, most network operators routinely rate limit the amount of ICMP echo traffic they allow. If bulk ICMP echo traffic cannot be managed, Internet links would become more prone to being targeted by packet flooding distributed denial of service (DoS) attacks.

Don't legislatively ban important network traffic management tools such as traffic prioritization and traffic rate limiting. All Internet traffic is NOT the same, and attempting to treat it as such may jeopardize the network's ability to support important applications such as voice and interactive video, or jeopardize the stability of the network itself, or its ability to manage attack traffic.

If, notwithstanding these issues, you do ban network traffic prioritization and traffic rate limiting, you may force sites to grossly overprovision/underutilize their network connections, thereby statistically ensuring that congestion virtually never occurs.

In other cases you will force sites to deploy multiple parallel connections (rather than a single shared/converged link) with each independent dedicated to a single type of traffic, thereby ensuring via crude manual separation of traffic that critical types of traffic can be protected, if only through through the use of dedicated links.

Either alternative will eliminate important technical options, lower efficiency, and increase costs to users.

II. For The Purpose of Any Network Neutrality Rules That May Be Adopted by the Commission, The Entities That Are Treated As "Internet Service Providers" Should Be Broadly Scoped/Inclusive

If the Commission *does* elect to adopt so-called "Network Neutrality" rules, those rules should be applied as ubiquitously and uniformly as possible to ALL "ISP-like" providers of Internet connectivity to end users.

That is, "provider of Internet connectivity" certainly include traditional commercial Internet Service Providers, but commercial ISPs are not the *only* ones who should be required to abide by any network neutrality rules the FCC may promulgate. The network neutrality rules should apply to any so-called "eyeball network" where network flows are dominated by inbound flows, from the Internet to diverse end users.

For example, when universities provide network connectivity for their faculty, staff and students, they effectively become indistinguishable from a "real" ISP to their faculty/staff/students, particularly when providing loosely-policed connectivity in non-academic settings such as on-campus housing. As such, university network operators are a classic "eyeball network," and one that should be held to the same network neutrality standards as traditional commercial ISPs.

As another example, when libraries (or restaurants) provide network connectivity for their patrons, they effectively manage their patrons' connectivity in a commercial-ISP-like way, and should also be subject to the same standards as a commercial ISP, whatever those standards end up being.

Does this mean that even corporate networks should be treated as being *de facto* ISPs? No. Unlike a university or library or restaurant (where end users typically have substantial discretion when it comes to how they use the network connectivity offered to them), corporate networks are normally provisioned for narrowly-defined purposes, such as sharing product literature with potential customers or processing customer orders. They are used for a specific purpose, not as end users may individually determine.

Another example of networks that should be exempt from network neutrality regulation are those that do not connect to "the Internet." For instance, a dedicated university network that is used solely to backhaul security video traffic from campus video cameras to a campus security office digital VCR should not be subject to network neutrality regulations since it doesn't "connect to the Internet."

In considering what networks do (or don't) "connect to the Internet," let me emphasize that Internet connectivity need not include exhaustive connectivity to "every" prefix or "every" autonomous system number. A network that connects to a "substantial fraction" of the Internet (or to at least major Internet properties such as Google or Microsoft), whether directly or through a gateway, should be considered as one that has "Internet connectivity."

III. Different Types Of Connectivity Are Not All Alike & Throughput is Not Always Within the ISP's Control

It has been suggested that all types of connectivity should be treated the same way -- wireline and wireless alike. This is another suggestion that I consider well-intentioned, but which fails to consider basic technical realities. We might like to pretend that all types of connections are "the same," but in fact they simply aren't.

For example, consider bandwidth. Fiber connectivity offers order of magnitude more bandwidth than WiFi or cellular data connectivity. Applications that are practical over a wireline connection (including things such as high bandwidth IP multicast video) may saturate wireless connections if we attempt to deploy them there.

Or consider the impact of bandwidth delay products. A remote user, far from the resource they're attempting to access, may see poor performance no matter how fast a local connection they buy simply due to poorly-tuned TCP/IP stacks. A novice user may visit a network "speed testing" site, notice that they only getting a fraction of their connectivity's bandwidth, and say, "Ah hah! My ISP is not giving me the bandwidth they owe me," when in fact the issue may not be a matter of their connection and its operation by the user's ISP, but a matter of how the user's own system has been configured or how it has *not* been optimized for maximum performance.

IV. Don't Prevent ISPs From Offering a Rich Portfolio of Market Options

Finally, I am concerned that well-intentioned network neutrality rules may inadvertently serve to limit the market options available to consumers. That is, some consumers may be cost conscious, and just want the cheapest option they can obtain, recognizing that compromises may be associated with how that service gets provisioned. For example, a cheap connection may have more users sharing a common uplink, leveraging the statistical reality that in most cases only a few users at any given time will be using that connection.

Others may want to buy a premium small business-class product offering tighter service level agreements or other advantageous characteristics, such as lower levels of uplink sharing, albeit at a higher price.

America has always been about options and choices, and I'd urge the Commission to ensure that the market can continue to have options and choices when it comes to its network connectivity. We don't need, and shouldn't end up with, a single "one size fits none" solution.

Thank you for the opportunity to offer comments on this matter.

Sincerely,

/s/

Joe St Sauver, Ph.D.