

Comments regarding GN Docket No. 09-191 and WC Docket No. 07-52.
by Jay Libove, CISSP

I have been an “Internet” user since the late ARPAnet days. I participated in a small way in the formation of one of the earliest and later largest and most successful, public Internet providers. I have been a commercial subscriber to Vonage VoIP service for more than five years, and had intermittently used early voice over Internet applications as much as five years before that. I have the technical and professional background to understand these debates, and the personal commercial experience as a user and customer of these services for the entire duration of their existence, to make informed commentary on what I want to see and pay for as a consumer.

My comments, where they reference specific paragraph numbers, take those paragraph numbers from FCC document 09-93A1 available on the Internet at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-09-93A1.pdf.

In codifying the principles of the *Internet Policy Statement* and converting them from user expectations to the specific obligations of identified parties, we must be sure to include all tiers of traffic carriers – ISPs, regional aggregation networks, and backbone providers, equally. The new statements of the principles refer to “a provider of broadband Internet access service”, which could be too-narrowly read as meaning end-user ISPs only. In particular, the definition in the proposed §8.3 of the “Internet” and of “Broadband Internet access” can be narrowly interpreted to mean the shortest path between a consumer user of Internet services and some middle part of the user’s ISP networks. This leaves unregulated by the proposed rules the rest of the path between the user and the service or content which he desires to reach. As requested by the FCC in the proposed rulemaking in paragraph 107, because individual carriers at the regional aggregation, interconnection, and backbone layers could also be swayed for economic reasons to non-transparently tamper with traffic priority and application source, the words used in the new statements of the principles must clearly also include these and all other distribution and backbone levels which make up the carriage of packets over the general Internet.

I am gratified to see that the FCC has remained with a conservative (in the sense of minimal regulation to achieve a goal) approach, and has written these proposed rules so that their applicability is clearly segregated to the businesses, and parts of businesses, which are the carriers of packets, and their applicability is not to include the businesses and parts of businesses which provide enhanced / value added / content specific information services, etc. This is similar to the regulatory wall which was imposed on the ILECs in the early days of DSL, to prevent anti-competitively preferential treatment of a phone company’s own DSL Internet service at the expense of the service which the phone company’s line department should provide to competitive DSL providers (DLECs). Extension of these principles beyond neutral packet carriage, to the actual providers of information services, would have too easily and most controversially put the FCC in to the very category of First Amendment risks which some commentators raise. It may be appropriate to discuss neutrality regulation of these information providers in a different proceeding, but separation should be kept in order to avoid these other issues from preventing the implementation of these most important basic net neutrality rules as presently proposed.

The FCC specifically requests comment as to whether the fourth principle (§8.11 Competitive Options) should be codified, or whether the first three rules (§8.5 Content, §8.7 Applications and Services, and §8.9 Devices) adequately achieve the intention of the fourth. While one could

reasonably draw the fourth principle's statements from the first three, one could also narrowly interpret the first three at the expense of the fourth. The first three are drawn in terms of the user. The fourth makes clear that providers of broadband Internet access services may not impinge on the user's rights by preferring or hindering any particular sources.

No information service provider, application provider, or content provider should be permitted to buy greater access to an ISP's subscribers. This would produce the same result as the ISP charging different prices for access to different information, or blocking some information from accessibility by its customers, simply with the change in money flow coming from a different place. The net effect would be the same: the consumer is prohibited from freely utilizing his purchased block of bandwidth for the services, sources, and content he most desires.

The fourth principle should be codified.

The proposed fifth principle (§8.13 Nondiscrimination), should be codified. Importantly, it must be drawn in a way so as to allow for user-chosen preferences, and for public standards driven good general network management practices, so as to avoid this principle becoming a roadblock to innovation. A good example of this is in the proposed rulemaking in paragraph 106's statement of the understanding of the term "nondiscriminatory". The text as proposed would fail to meet this broader flexibility, which I as a user of latency and jitter sensitive VoIP services would like my ISP to be able to offer to me, transparently and at my option.

Additionally, as some comments have suggested, it may be appropriate to clarify this principle in terms of anti-competitive behaviors. This may possibly be achieved by adding a few words to §8.13 thus: "Subject to reasonable network management, a provider of broadband Internet access service must treat lawful content, applications, and services in a nondiscriminatory manner, with specific attention paid to discrimination which could be construed as anti-competitive."

User permission or preference for giving priority to certain traffic may be signaled either explicitly through a user's selections in a control panel in their Internet Service Provider user account, and/or implicitly by a user's settings in their customer controlled routers and software configurations at their customer premises. An obvious example of this would be for a consumer with VoIP software or hardware to express a preference for the VoIP traffic to have priority over most other traffic to/from their endpoint. It would not be a violation of net neutrality principles for ISPs and backbone providers to make reasonable efforts within their own traffic shaping and network management infrastructures to honor these users' preferences, within each user's own purchased bandwidth allotment.

To the degree that any types of traffic are given priority, by the independent choice of an ISP or backbone provider, absent the explicit request (through user control panel settings) or implicit request (through software and hardware, under the control of the consumer, marking traffic to indicate those preferences), such priorities should be broadly and transparently drawn based on standards and proposed standards relating to such traffic prioritization. A standard might be interpreted to suggest that a protocol should be treated as high priority if the standard clearly addresses the negative or even preventive impact that high latency and/or jitter would have on the successful use of the described protocol. For example, highly interactive forms of traffic should be broadly equated for placement in to high priority traffic categories (VoIP, interactive gaming, interactive remote session control applications such as telnet, SSH, remote desktop, etc). This reliance on standards and public discussions of proposed standards, transparently enacted by providers, prevents the need or tendency of regulation to become too specific, at the expense of future innovation, by making transparently available to the interested public the details of how this type of network management practice is carried out. Should a new application begin to gain

popularity, and suffer by not being granted membership in the generally agreed class of high priority traffic applications, there could be no complaints of dark practices by ISPs, because the practices are transparent. An open discussion would then ensue, coming to a public agreement through accepted standards processes as to whether the new application should belong to the general high priority class, or should remain within general traffic. Similarly, there may be broad public agreements as to lower priority (bulk) traffic. All of this is conditioned upon these priorities being implemented within each user's own purchased bandwidth block. It is in the general interest of the consumer subscriber, and within his understanding, to agree that when his total bandwidth use exceeds his purchased bandwidth block, he may either reduce his own use, or increase the amount he is willing to pay to have more bandwidth allocated to him. Finally, because time / jitter sensitive traffic – generally speaking the kinds of traffic which would reasonably be proposed to benefit from placement in high priority groups – tend to be relatively low total bandwidth compared to the kinds of bulk transfers (especially peer-to-peer programs) which are the primary cause of congestion concerns, there should rarely be a point where it can be claimed to be unreasonable to request addition of a new sensitive application to the higher priority tier, on the basis that adding yet another high priority protocol would overwhelm overall capacity, as broadly agreed lower priority traffic groups should always be able to lose that small additional amount of bandwidth reserved to the higher priority traffic group without major user impact. (This relates to the request by the FCC in paragraph 137).

If total ISP uplink or backbone capacity becomes saturated to the point where reasonable network management practices and the above mentioned publicly agreed, transparent, and consumer accepted traffic priority classes can no longer adequately manage total traffic, and all traffic begins to suffer, then probably again we reach the point where the consumer subscriber will understand and agree to either back off of his less preferred traffic, or increase his total payment to allow higher total amounts of traffic to be carried Internet wide. Most likely, this possibility will never be realized, as reasonable bandwidth price tiers should support adequate overall backbone capacities long before total backbone saturation significantly degrades total performance.

Regarding these neutrality principles vs. the concept of prohibiting only “unreasonable discrimination”, I propose that the “unreasonable discrimination” standard is too narrow, and would embolden unscrupulous operators to push the lines too far. The neutrality principles make clear that there will be limits, and that public discourse and educating of the regulators in advance – in other words, cooperative, collaborative, and transparent practices – are the method more likely to meet with commercial success with lower risk of regulatory and legal reaction. The concern that the more general neutrality principles would prohibit “socially beneficial discrimination” which the more narrow “unreasonable discrimination” standard would not, is to be met with transparency. Even though the neutrality principles are to be written as rules, the rules – particularly §8.13 regarding Nondiscrimination – should be written with sufficient flexibility for experiments in socially beneficial discrimination to be conducted transparently without fear of regulatory over-reaction.

The proposed sixth principle, §8.15 Transparency, should be codified. In barely competitive or uncompetitive markets (the majority of the United States), there exists an imbalance of power between the commercial provider and the consumer. In such situations, regulatory requirements for disclosure are needed. This sixth principle must be interpreted and enforced in a way to prevent intrusion in to competitive advantage as to the efficiency and secrecy of the exact methods by which network management are accomplished, while assuring that the interested party – ranging from the non-technical consumer through the highly specialized developer or service provider will have ready and comprehensible access to information sufficient to understand how their intended user or provision of information services will be affected by the Internet access provider's “network management” practices. These disclosures must, at a minimum, take place outside of the detailed

and dense legal text of consumer service agreements, for example in well organized FAQ documents available to subscribers and non-subscribers alike.

Specific types of disclosures to users might include the percentage of oversubscription of the Internet access provider's uplinks, along with graphs showing whether and when this oversubscription ever actually results in user-observable degradation in performance, with examples of how the user might experience that degradation. Oversubscription is an appropriate practice, and when well managed will rarely result in too much user-observable degradation at too-bad a level. Well structured FAQs and demonstrations should convince users of this. Another type of disclosure should be showing the tiered priority classes, their saturation, and their relationship to total bandwidth used and to total bandwidth available. This disclosure should answer for the user his doubts as to whether his VoIP or gaming or other sensitive information service is not working well because of something occurring within the control of his Internet access provider, or beyond that control. This should even have the beneficial side effect of reducing customer support contacts to Internet access providers for issues which are beyond the control of the provider. These types of disclosure will produce – where competitive choice within a geography is even an option at all – a more competitive marketplace, by permitting the user to see factual information about what does, and what does not, affect his Internet experience, and whether the competitive provider would do any different. These disclosures should remain available as a historical archive, to allow users, public interest groups, and regulators to see trends, compare current performance to past performance, and watch out for revisionism in FAQs and current disclosures by providers.

When a provider plans to make changes to its network management practices which could result in calculable or provable differences in user experience (even if the great majority of users would be unlikely to actually notice the difference), these changes must be clearly notified to the users and the public in advance of their taking effect. As above, this could have the beneficial effect of a) reducing customer support contacts for which “nothing is actually wrong”, and b) enhancing the value of customer contacts where the provider's estimation of customer impact or the degree to which the impact would be noticeable by the customer turned out to be incorrect, by enabling the customer to say “I notice X change in my service today, and that you announced a network management change yesterday, and they seem to be related” rather than “my service isn't working as well today as it did yesterday, what's wrong?”.

Another disclosure trigger – which an Internet access provider should be doing already today for simple reasons of good customer service and to manage customer support contacts – should be when the Internet access provider takes a network management action, e.g. in response to a denial of service attack, or a worm, which may have undesirable user-observable side-effects, but which should be relatively short lived until the emergent situation (the attack or worm) subsides, and which on the whole are to the benefit of the group of subscribers to the provider. Such reactions, unless egregious, should generally have a presumption of innocence against claims of provider violation of neutrality principles.

To the degree that a provider feels that the quantity of information necessary to meet a disclosure requirement exceeds the degree to which it could make a disclosure without violating its competitive need for secrecy, certain narrowly limited parts of the disclosures could be made under seal to regulators, rather than in its entirety to the public. The balance of the disclosure, absent only those narrowly limited parts which would violate the provider's competitive need for secrecy, must still be broadly publicly disclosed, possibly with an indication of where redactions / private disclosures to regulators have been made.

To avoid concerns of “over-regulation”, or too-high burdens from disclosure requirements relative to the benefit the disclosure is intended to produce, disclosure requirements should start relatively modest, with much input from the regulated industries as to what their existing management systems make readily available, vs. what other types of disclosure might be quite burdensome to produce.

Regarding the interplay between these general disclosure requirements, which take place over the lifetime of a subscriber’s relationship with a provider, and the law enforcement exceptions, which generally take place for short periods of time and generally should not have large user-observable impact, the exemptions codified for law enforcement purposes should not have a negative interaction with the disclosure requirements, and the disclosure requirements should not have to take in to account law enforcement action, but providers should be cautioned to avoid using the law enforcement exception to skirt disclosure requirements in any way.

Personal privacy (paragraph 130) should not come in to play at any level of aggregation above the display of an individual subscriber’s own bandwidth utilization chart. Such privacy concerns should be handled by a clear statement of the personal responsibilities of the subscriber to inform his family, guests, etc that such information will be disclosed by the provider to the subscriber, and as a result certain categories of use by the subscriber’s family, guests, etc may come to the subscriber’s attention. However, beyond this notice, there should be no personal privacy concerns in the operation of net neutral services.

It is not specific to neutrality rules, but general to the trusted position of an Internet access provider, that privacy impact assessments should be performed on all service offerings, and where personal privacy could be violated, appropriate procedural, technical and physical controls should be implemented, and appropriate disclosures made, to balance those privacy risks with the value of the service provided to the consumer. This should be a separate discussion from net neutrality.

The proposed “specialized services” (IV.G), which are in fact services provided over an Internet type of infrastructure, but which are not themselves Internet access services, as the FCC proposes should not be subject to these principles; but in return they must be subject to the same regulations to which all competing services are subject. In other words, a telephony service is a regulated telephony service, regardless of how it is carried. Vonage and Skype are Internet applications whose packets when carried over the Internet should benefit from these neutrality rules, but AT&T telephony carried over IP protocols within separately managed, reserved bandwidth is not an Internet application and is not affected by these neutrality rules. The non-Internet dependent parts of Vonage and Skype service may be subject to the Telephony regulations, in so far as they are classified as falling in to the category of regulated Telephony services. Net neutrality principles would apply to the data packets of Vonage and Skype carried over the Internet, but not to the AT&T telephony service which travels over reserved bandwidth which is operationally, and in the user experience, unrelated to their general Internet bandwidth. Similarly for a Cable TV service. AT&T’s U-verse “Cable TV” (over IP) service, in so far as it is sold as Cable TV, and only incidentally runs over the same cables over which AT&T also sells Internet access, but as a completely separately managed service with its own different bandwidth allocation, is a Cable TV service and should be regulated as such, not as an Internet access service.

Note that this distinction must fail if the total bandwidth in the pipe sold as an Internet access service to the consumer does not have a hard partition between that part used for Internet access and that part used to deliver Cable TV or Telephony. In other words, if the user’s act of turning on a second, or third or fourth “television” or “telephone” in the house would reduce his available

Internet access bandwidth, then this regulatory wall between the Internet access part of that subscriber's connection and the Cable TV part of that connection would have to be weakened, as the act of devoting further bandwidth to the Cable TV service would unavoidably result in less bandwidth available to another member of that subscriber's household streaming e.g. a YouTube video. This answers one of the FCC's questions in paragraphs 150 and 151, regarding the technical characteristics which may put a service in the specialized category, or prevent a service from benefitting from that category.

Examining a service from the other side, where a company which historically has provided a regulated content service is now offering limited access to that service by forwarding it over the general Internet, the recent moves by some Cable TV providers to grant access to subscribers' subscribed content via the Internet ceases being a separately regulated Cable TV service in so far as it impacts or is impacted by the net neutrality principles. When a cable TV subscriber logs on to their cable TV provider's website, authenticates himself to gain access to his subscribed content, and then streams that content over any general Internet access service (no matter that the service may be provided by the same company, or company umbrella, as their cable TV provider; or may be a completely unaffiliated Internet access provider at the hotel where he is presently staying), the net neutrality principles must fully apply, and the delivery of the bits of the streamed media from that authenticated cable TV over Internet service must be treated equally with e.g. a YouTube stream.

Generally speaking, because managed or specialized services rely on provision over a specific provider's infrastructure, whereas general Internet-carried information services (such as Vonage and Skype) are intended to work over any Internet infrastructure, competition between specialized services and services intended to work over the general Internet is somewhat constrained. In particular, a person who relies on Skype for his phone service – home, office, travelling – is not competing that Skype service against his Internet provider(s)' Telephony (over IP) services because his provider(s)' services are not usable in the variety of places where Skype is usable. This, to a small degree, limits the concern stated by the FCC in paragraph 149 that a too-broad exemption from the neutrality rules for specialized services may ultimately result in less competitive protection than desired from these rules. Still, the FCC is correct to be cautious in this area, and this rulemaking might address this concern through a commitment to monitor and report on those services which come to fall under the managed or specialized category over time, as to whether those services produce the kind of unfair competition to services over the general Internet which these neutrality rules are intended to protect and promote.

Regarding tethering (paragraph 167), I propose that a basic question to ask, before deciding whether to regulate tethering, is: Is there any difference to the wireless carrier's data network between a laptop PC tethered with a data-capable mobile phone device, and that same laptop PC connected with a "3G modem" (USB or PCMCIA), other than the total amount of data which a laptop PC may be expected to consume as compared to the presumably lesser amount of data which the data-capable mobile phone device may be expected to consume? If the answer, as I suspect, is "no", then I propose that contractual limitations on the use of data-capable mobile phone devices for tethering purposes should be prohibited, and the same bandwidth pricing tiers be used to effect commercial control on the cost and service impact of wireless data.

The FCC requests comments on the "digital divide". I do not see a connection between the overwhelming socially beneficial desire to enfranchise those socio-economic groups which today participate at significantly lower rates in broadband Internet access (the "digital divide"), and net

neutrality. I believe these two issues should be discussed separately.

Speaking philosophically about some of the public debate which has raged in recent months, “net neutrality” has been debated heavily, often with extreme rhetoric which we can only conclude is intended to frighten consumers, businesses, and the FCC away from any form of regulation. Such extreme claims have been made as that these proposed conservative and basic neutrality rules are the camel’s nose under the tent, which will result in the FCC eventually regulating the details of speech content on the Internet, in violation of the First Amendment. This might be a valid concern if the FCC were proposing to regulate what could be controlled, on a content basis, by ISPs or backbone providers. However, net neutrality as presently proposed simply and non-preferentially prohibits the control by ISPs or backbone providers of traffic based on its content. Therefore, these First Amendment arguments against net neutrality are invalid.

Any burdens placed by neutrality principles on Internet access providers’ right to (commercial) free speech under their status as “legal persons” for purposes of the First Amendment must be subservient to the more important right of users as natural persons to exercise their First Amendment guarantees.

Some opponents of net neutrality regulation argue that “some traffic imposes greater burdens on the network than other traffic and that “innovation could be even better for consumers if it could respond to price signals from platform providers,” such as by “tak[ing] into account potential congestion costs of bandwidth-intensive applications”. Indeed, a larger quantity of traffic imposes a larger burden on a network than does a smaller quantity of traffic. However, the content – source, destination, port, protocol, application – of that traffic has no bearing on the burden the traffic places on the network. Therefore, it does not follow to oppose net neutrality as it has been proposed by the FCC, on the basis of “different” traffic imposing different burdens on the network. And the proposed net neutrality rules in no way block ISPs and backbone providers’ right to use reasonable network management practices to control the negative impacts of high bandwidth utilization by individual subscribers. Therefore, these arguments against net neutrality are invalid.

I thank the FCC for its detailed and well reasoned attention to these potential issues.

Jay Libove, CISSP, CIPP
Atlanta, GA and Barcelona, Spain