Before the

Federal Communications Commission

Washington, D.C. 20554

In the Matter of the Open Internet Order Remand

GN Docket No. 14-28

April 2, 2014

Comments of NetAccess Futures

NetAccess Futures, through its Principal, Daniel B. Grossman, hereby submits these comments in response to the notice dated February 19, 2014 regarding the guidance of the United States Court of Appeals for the District of Columbia Circuit in the Verizon v. FCC case.

I. INTRODUCTION AND OBSERVATIONS

A. About NetAccess Futures

NetAccess Futures ("we") is a boutique consultancy, specializing in broadband access networking technology and strategy. Its Principal, Daniel B. Grossman, has over 34 years of experience in data networking technology. He has been deeply interested in the interaction between communications technology, economics and policy. Past experience also exposed him to the Commission's Rules and procedures and valuable role in maintaining competitive markets.

We have not engaged counsel to prepare these comments, which therefore no doubt deviate from normative form. As we have had to submit them late, we have also not footnoted as extensively as we might, and some footnotes are incomplete. Further, we have not developed all of our arguments as fully as we might have preferred. At this late date, we think it more valuable to have our observations and recommendations in the record, and hopefully of use to the Commission.

NetAccess Futures does not presently have any clients who are parties in this matter.

B. Background

The set of issues now referred to as "Open Internet" have a long and painful history, elements of which predate the Internet as we know it. This background is recited in detail in the DC Circuit opinion in *Verizon v. FCC*[1]. In the interest of time, space, and the readers' patience, we do not reiterate this history.

C. The Commission's Situation

The DC Circuit's ruling in *Verizon vs FCC* leaves the Commission in a fine mess. The Commission recognizes that certain hypothetical operator practices are not in the public interest. Its decision to classify Internet services as "Information Services" under Title I of the Telecommunications Act of 1996 was found by the court to restrict its authority to regulate such practices, as such regulation would be "Common Carrier-like". However, the Court did leave the Commission a narrow ledge to stand on, by finding an affirmative grant of authority in Section 706(b) of the Act.

The Court's decision has caused a renewed storm of public outrage, fanned by public interest groups and intensified by resurgence of economic populism. Thus, the Commission must do something. Yet there is every reason to believe that any new Open Internet Order will be appealed yet again by a broadband provider.

Many commenters prescribe Reclassification of broadband Internet providers under Title II of the Act. We are somewhat sympathetic to that idea. However, as a practical matter, that horse has bolted. Reclassification would be extremely difficult to accomplish, would likely attract negative attention from Congress, and almost certainly would be appealed. If the Commission were to reclassify, it would then be faced with having to conduct a formidable number of forbearance proceedings in order to suspend Rules that are inappropriate to the broadband Internet. We therefore appreciate the Chairman's desire to find alternatives.

D. How the Commission can break the cycle

We observe that the issues at hand are a giant hairball. They conflate many loosely related concerns. They are often excessively abstract and raised in terms of hypotheticals. They are grounded in distrust,

[1] cite

some of which is founded. Thus, the Commission can best proceed by decomposing the problem into manageable parts.

The positions of the parties often betray misunderstanding of the underlying technology and economics of the Internet. They are based, for example, in over-simplifications, false analogies, obsolete assumptions, and lack of grounding in core principles. We offer a small part of the subject matter expertise to correct the record.

It is not clear that the parties have a common understanding of the issues in dispute, and thus seem to argue past each other. Their claims are grounded in conflicting narratives, and as a result, each party's arguments can only be understood in context of their own narrative.

The parties make conflicting and anodyne claims about "freedom", "innovation", "investment", "free markets", "jobs", "competition" and so on. Such claims, while appealing in the abstract to important values, do not add any substance and cover for the parties' naked self-interest. The Commission should not be distracted by them, and instead focus on concrete value and harm.

II. "NET NEUTRALITY" IS A PROXY FOR OTHER ISSUES

The Open Internet Rules were intended to protect consumers from a closed set of abusive practices. However, the furor surrounding them has made them symbolic of a broader set of concerns. The Commission would be wise to sever these extraneous issues from this docket.

A. Peering, Transit and Settlements

Infrastructure-based communications providers, intermediaries and content providers have engaged in a long-standing tussle over their respective shares of revenue extracted from consumers. "Network Neutrality" and its association to the public interest is largely a cover for this tussle. Indeed, the first shot in the "Net Neutrality" wars was fired by former SBC Communications CEO Ed Whitacre in an interview with Business Week[2]:

> "…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 this 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. Why should they be allowed to use my pipes?"

---

[2] "At SBC, It's All About Scope and Scale " Business Week, Oct. 24, 2005

Ultimately, consumers (or advertisers as proxy for consumers) bear the entire cost of the value chain from content creation to browser. Broadband providers are paid by consumers for connectivity to the Internet, almost universally at a flat rate. Content providers are paid for content directly by consumers, and/or by advertising. They, in turn, pay a wholesale provider for connectivity, or connect directly to the broadband provider. Broadband providers claim this system compensates them insufficiently to recover necessary capital investment in their networks.

Presently, the flash point of the battle over revenue allocation is at the point of interconnection to the broadband service provider's network, e.g., in the recent spats between Netflix and Comcast[3], and between Level3 Communications and Comcast[4]. This problem has its roots in the way interconnection works in the Internet[5]. Historically, it was assumed that when two providers interconnect to provide paths between their own subscribers, the flows of traffic between providers are roughly balanced. The networks are thus "peers", benefitting equally from the interconnection. Networks that provide paths between two or more other networks are called "transit networks". Peering arrangements are usually settlement-free; however, some include penalties for traffic imbalance. External costs of interconnection are typically shared. Often, peering agreements are informal, not subject to contract, and terms and conditions of peering are typically confidential[6]. Summaries of peering policy may be published in a Peering Database[7], and networks may publish their peering policies. However, any disclosure is entirely voluntary, and usually limited.

As unidirectional video traffic has come to dominate Internet utilization, the assumption of traffic balance ceased to hold. Thus, economic benefit of interconnection accrues to content providers and intermediaries, while broadband providers disproportionately bear capital cost of incremental capacity to support increased traffic volumes. When broadband providers refuse to add capacity and/or demand payment for imbalance, content providers or intermediaries raise a hue and cry about "violation of net neutrality". Consumers are collateral damage in the resulting battle.

[3] e.g., see "Comcast and Netflix Reach Deal on Service", The New York Times, February 23 2004
[4] e.g., see "Netflix Partner Says Comcast 'Toll' Threatens Online Video Delivery", November 29, 2010
[5] e.g., see Brough Turner, "Impact of Internet Peering on Network Architectures and Economics" OFC 2014, also at http://blogs.broughturner.com/2014/03/internet-peering-an-invited-talk-at-ofc-2014.html
[6] Woodcock, Bill and Vijay Adhikari, "Survey of Characteristics of Internet Carrier Interconnection Agreements," Packet Clearing House, May 2, 2011, https://www.pch.net/resources/papers/peering-survey/PCH-Peering-Survey-2011.pdf
[7] www.peeringdb.com

While existing interconnection arrangements between ISPs are apparently mostly benign, the presently unregulated status of interconnection is becoming a problem. Furthermore, since peering terms and conditions are typically confidential, they leave room for unfair discrimination.

The Commission should therefore open an Inquiry to consider what Rules, if any, ought to apply to interconnection arrangements and settlements in the Internet. In particular, the Commission should consider whether the opaque nature of peering and transit agreements causes harmful market distortions. All matters related to interconnection and compensation may then be redirected from the Open Internet proceeding.

B. Competition and Regulation in Broadband Access Markets

Much commentary[8] has concerned alleged abuse of monopoly power by broadband providers. On the other hand, parties on the other side[9] claim plentiful competition. Both sides refuse to acknowledge the special economics of broadband access markets. New entrants are deterred by the exceptionally high cost of capital infrastructure, where most of this cost is not incurred as a direct result of sales. Further, market dilution and incumbent advantage reduce expected return on investment. The business case for being a second entrant diminishes further for the third entrant, yet further for the fourth, and so on. In short, hope that competition can resolve the "Open Internet" problem is not realistic.

Parties also raise the specter of excessive regulation[10]. This too is risible. Absent vibrant competition to enforce market discipline, and absent vigorous regulation, consumers will inevitably suffer under oppressive monopolies in the Broadband access market.

We observe that access markets, because of their low density and price sensitivity, are far less amenable to competition then other market segments in the Internet's ecosystem. The Commission should be mindful of this distinction, and recognize that it must result in differing rules.

C. Consumer Dissatisfaction with Incumbent Broadband Operators

Many consumers and pro-consumer interest groups conflate concern for the Open Internet with unrelated complaints about broadband providers. Various commentators to Internet fora[11] on the

[8] e.g., cite, cite
[9] e.g., cite
[10] e.g., cite

subject of the Open Internet proceedings offer numerous complaints: excessive prices, poor customer service, reliability issues, performance issues, privacy concerns, failure to invest in performance upgrades, failure to extend high-speed broadband services into high-cost areas. Opposition by broadband providers to the Open Internet order (and its predecessors) is seen as part of a pattern of abusive and anti-competitive practices. Indeed, broadband providers are often cast as the scheming, greedy, evil pasteboard villains in a morality play.

The Commission cannot make progress on the Open Internet unless it is able to sever these concerns, and deal with them in proper course. This may require new rulemaking proceedings, enforcement actions, merger conditions, public information statements, or other actions by the Commission. We can offer no specific prescription or roadmap.

III. BLOCKING SHOULD BE ADDRESSED AS TWO SEPARATE ISSUES

We turn next to the issues raised by the Open Internet Order. It prohibited fixed broadband providers from "block[ing] lawful content, applications, services, or non-harmful devices, subject to reasonable network management" and forbade mobile providers from "block[ing]consumers from accessing lawful websites" and from "block[ing] applications that compete with the provider's voice or video telephony services, subject to reasonable network management."[12] We observe that two distinct potential harms are in play here, and can be more easily resolved separately than together.

A. Blocking motivated by point-of-view should be taken off the table

Interest groups and consumers are concerned that broadband providers could block access to content or services which express a point-of-view that they disagree with. For example, Free Press frets:

> "… a company like AT&T or Verizon could decide where their users can go for news and what stories get buried or blocked online. Verizon could strike a deal with CNN and hamper their users' ability to access alternative news sources. Comcast could slow access to Al Jazeera, because it wants to promote its NBC news offerings."[13]

[11] e.g., see reader comments on "Will the Net Stay Neutral?", www.nytimes.com, February 21, 2014
[12] cite
[13] Josh Sterns, "Net Neutrality and the Future of Journalism", http://www.freepress.net/blog/2014/01/15/net-neutrality-and-future-journalism

Such a scenario, were it to play out, would be indisputably be deeply offensive to American values of free expression and open political discourse. Of all the parade of the horribles raised in the Open Internet debate, this one could be seen as the most disturbing.

However, capability does not necessarily mean intent. It is difficult to imagine a rational business motivation for such an action. Even if they wished to support or oppose a political point of view directly affecting their business interests – e.g., by blocking the Free Press website – a rational business person would weigh limited benefits against risks of public opprobrium, enforcement action by the Commission, and litigation.

One could imagine a small, privately held broadband provider that acted in a less rational way; e.g., one that chose to block lawful web sites deemed “immoral” by their owners. For this reason alone, there is practical value to maintaining having a rule against blocking motivated by point-of-view.

A possible way out seems worth exploring.   We expect that the largest Broadband providers would consider point-of-view blocking to be unthinkable. If this is the case, the Commission could solicit permanent and binding commitments from the broadband providers that they would not engage in such practices.   This commitment could be followed by a negotiated Rulemaking, narrowly tailored to prohibiting point-of-view blocking.

B. Blocking with anti-competitive motives

Broadband providers might also be motivated to block services or applications that compete with similar services or applications provided by an affiliated entity.   Such actions have precedent. For example, in the *Madison River*[14] case, a wireline provider blocked competing VoIP services in order to protect its own voice service.   The Commission used its enforcement authority under Section 201(b), ultimately obtaining a consent decree prohibiting anti-competitive blocking. However, this depended upon Madison River’s status as a Common Carrier, which would not apply to broadband providers under the current classification.

The Commission must, instead rely on its authority under Section 706(a), and its authority to refer anti-competitive practices to the Federal Trade Commission and/or the Anti-Trust Division. These can be potent deterrents. As this is a purely legal matter, we can offer no prescription as to how this might be implemented.

---

[14] DA 05-543

IV. "UNREASONABLE" TRAFFIC DISCRIMINATION SHOULD BE NARROWLY DEFINED

The Open Internet order prohibited fixed broadband providers from "unreasonably discriminat[ing] in transmitting lawful network traffic"[15]. This, arguably, is the most problematic and controversial element of the Open Internet Order. We suggest that the Commission re-examine its reasoning on this subject[16] to more narrowly tailor this prohibition to address concrete competitive harms, while permitting services and network behaviors which are beneficial to consumers.

We here must digress into a technical explanation of Internet technology. Some of our points are simplified for clarity.

A. Why different classes of content require different treatment

There are numerous important distinctions between content which is transferred across the Internet as unitary data objects, content which is transferred as a stream of bytes, content which takes the form of short, independent messages and content which synchronizes the states of systems with independent inputs. Examples of the first category include text, photographs and animations, programs, web pages, documents and files. Examples of the second includes streamed video (including television), streamed audio (e..g, as in Internet radio), telephony and video conferencing. Examples of the third include home security and energy management systems, and home medical devices. Examples of the fourth include many kinds of live action games[17]. Each of these categories benefits from different treatment by the network in order to maximize end-user satisfaction.

We introduce here the notion of a "data flow"[18]. A data flow, in this context, is a sequence of packets from a source to one (or, in the case of multicast) more destinations. It is the logical equivalent of a connection. All packets in a data flow have a common source address, destination address, port number, and protocol; this property is used to associate packets with data flows. Data flows are unidirectional.

---

[15] cite
[16] Open Internet Order, Part III C (2) at 68--79
[17] In the interest of time and space, we will not elaborate further on the latter two categories
[18] The term "microflow" is used interchangeably. See RFC2475 at page 5. Internet Requests for Comment (RFCs) are obtained at www.rfc-editor.org. In the interest of time and space, we refer to them only by number.

i. Object Transfer

Consider a digital photograph. It is stored on a server and requested by a user through a web browser. Upon receiving a request, the server retrieves the photograph, segments into into multiple packets and sends those packets sequentially across the Internet, to the user's browser, which reassembles and displays them. The determinants of the user's satisfaction with this transaction are i) that the photograph is completely received and correctly displayed and ii) that the latency from request to display (or, colloquially, "click to view") is subjectively acceptable. The latter criterion is somewhat fuzzy. Its upper bound is the limit of the viewer's patience, typically a few seconds. Its lower bound is the limit of human perception, on the order of hundreds of milliseconds; human brains cannot distinguish, for example, between 50 ms and 100 ms of latency.

Objects like photographs on the Internet are usually transported using the TCP[19] protocol. TCP is characterized by a complex dynamic behavior[20], which is used to regulate traffic flow in the Internet. We describe this further below.

Data flows consisting of objects are characterized as "elastic"; they adapt their rate to network capacity[21] available at a given time, under the closed loop control system inherent in TCP. They are also characterized as "non real-time"; they are insensitive (within limits) to the timeliness of delivery of any particular packet.

ii. Stream transfer

Next, consider an encoded high-quality digital video. It is stored in on a server and requested by a user through a browser with embedded media player. Upon receiving the request, and the server retrieves the video and proceeds to send it across the Internet as a sequence of packets. The rate at which it sends is consistent with the video's "frame rate" (typically 30 or 60 frames per second), motion or changes in the scene between successive frames, and numerous other factors. Note that one frame does not necessarily correspond to one packet. The media player stores the initial packets of the stream in an elastic, first-in-first-out (FIFO) buffer until a particular number of packets have been stored. It then reads from the buffer into a decoder, so as to be able to display video at the frame rate.

[19] see RFC793; also note RFC2914
[20] see RFC3918 for a summary of TCP dynamic behavior.
[21] "bandwidth" is often used interchangeably with "capacity", e.g., of a link, in bits per second; for purely pedantic reasons, we choose the latter

In order to maintain satisfactory video quality, the packets required to display the next video frame in sequence must always be available in the elastic buffer at the instant the decoder needs them. Otherwise, the video will, for example, appear to pixelate, skip, or freeze, or the player might lock up. These detracts from the user's experience.

Observe that the elastic buffer is filled and drained by asynchronous, decoupled processes. Assume, for simplicity, that the server emits packets at a constant rate of R packets per second, at intervals of 1/R seconds. As these packets cross the Internet, they are subject to buffering in each switch or router that they encounter along the path to the user. The depth of any switch's buffer at a given time varies dynamically according to the aggregate behavior of all the data flows that share that buffer. As a result, the delay experienced by each packet, and thus the inter-arrival time between packets at the elastic buffer varies. The decoder, however, drains packets from the elastic buffer rate R and interval 1/R. This is the point of the elastic buffer: it absorbs delay variation to assure that the decoder will have the next packet available when needed. At the beginning of a video, the elastic buffer must be filled to a pre-determined depth. A very deep elastic buffer will noticeably delay the start of play out. A too-shallow elastic buffer will become exhausted in periods of congestion. In addition, elastic buffer depth is typically limited by hardware constraints. Thus, it is highly desirable that delay variation be bounded.

Streaming data flows are characterized as "inelastic"; they operate at characteristic rate which is determined by the application and/or content, and do not continuously adapt to varying network capacity[22]. They are also characterized as "real time": they are sensitive to timeliness of delivery of individual packets, to the extent that a late-arriving packet is as good as a lost.

iii. The canonical Internet offers "Best Effort" service

We now turn to the behavior of the canonical Internet.

The Internet provides a "Best Effort" service[23]. This implies the following:

a. the data flow does not have a characteristic rate

---

[22] This discussion concerns high quality video. Various schemes for "adaptive bit rate" (ABR), variable quality video are also widely used. These "gear shift" between stored segments of the video of progressively greater rate (and thus better quality), so as to roughly match available network capacity over each segment.

[23] Cite

b. the source of a data flow may transmit packets freely, without any enforced time or rate constraint

c. each packet is temporarily stored in a buffer at each successive hop (i.e., switch or router) along the path of the flow. Packet buffers organized into FIFO queues. These queues are shared by multiple flows which have a common next hop. Packets are transmitted in the order in which they are received, regardless of which flow they belong to

d. packets wait at each hop for a variable amount of time

e. the depth of each queue is monitored, and one or more packets are discarded if it exceeds a pre-determined threshold

f. variable delay and packet discarding are entirely normal and expected behaviors, and are implicit signals that at least one link along the path is congested

g. if every data flow, in response to congestion signals, adjusts its rate as required by TCP, the rates all flows will converge to an approximately "fair" share of each link along their respective paths

h. data flows that do not conform to normative TCP behavior might utilize an "unfair" share of link capacity, potentially increasing variable delay and/or packet drop rate for all flows

Best effort service is well matched to elastic applications. In reasonably well engineered networks, the rate of any data flow varies in the presence of competing traffic, yet this variation yields nominally acceptable latency.

iv. "Best Effort" can be augmented by strict precedence in the Internet

Contrary to commonly held belief[24], the Internet Protocol (IP) has always supported a precedence mechanism[25]. This mimics the National Communications System's multi-level precedence and preemption system. In addition, precedence is used to assure timely delivery to critical network control traffic (such as certain kinds of routing updates). The IETF's "Requirements for IP Version 4 Routers"[26] states:

---

[24] cite needed

[25] see RFC 791, "Internet Protocol", September 1981

[26] RFC 1812, June, 1995

> "Routers SHOULD[27] implement precedence-ordered queue service. Precedence-ordered queue service means that when a packet is selected for output on a (logical) link, the packet of highest precedence that has been queued for that link is sent. Routers that implement precedence-ordered queue service MUST also have a configuration option to suppress precedence-ordered queue service in the Internet Layer.

Precedence is useful in tactical and public safety networks, where strict policies and procedures govern its use. In less-structured environments like the public Internet, and absent pricing disincentives, users are incented to use the highest precedence at all times. This renders the highest precedence meaningless, and thus becomes a Tragedy of the Commons. Therefore, precedence-ordered queue service is typically disabled in the public Internet. If a user sends a packet marked with other than the "Routine" precedence, the network returns an error message.

v. The Internet Protocol has long supported services other than Best Effort

Research in the Internet community on streaming audio and video dates to 1979.[28] We will not recite the history of the early experimental work, in the interest of brevity.

In 1994, the IETF published the first of a body of work called "IP Integrated Services", or "Intserv"[29]. Intserv extends the Internet architecture to support "real-time services" (Best Effort is considered "non-real time"). The Intserv effort produced two subtly different services: "Guaranteed Service" and "Controlled Load Service". These services are applied, on demand, to individual data flows. Intserv was widely implemented in routers and in computer operating systems. However, it was found not to "scale" well, due to the amount of persistent state that routers need to retain, and the fact that all routers along a path must support it if it is to work at all. Thus, it did not achieve widespread use in the Internet.

Intserv was followed in 1998 by a somewhat different architecture called "IP Differentiated Services" or "Diffserv". Diffserv specifies a "tool kit" from which network operators can create services. It is designed to perform processing- and memory intensive functions only in the routers at the edges of the network. Edge routers "classify" packets based on their overhead information, and "mark" them according to a policy (in addition to other performing functions). Routers in the core realize relatively

[27] "SHOULD" "means that "there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course."
[28] James Forgie, IEN 119, "ST – A proposed Internet streaming protocol", September, 1979
[29] see Braden et al, "Integrated Services in the Internet Architecture: an Overview", RFC 1633

simple "per-hop behaviors", which direct packet handling based on each packet's marking. Diffserv is widely implemented in routers, and commonly used. In particular, broadband providers offering IP telephony, IP video, and enterprise services use Diffserv in their "managed" network[30]. Technical capabilities for offering Diffserv-based services to residential broadband subscribers are implemented in cable[31], DSL[32] and fiber-to-the-home systems[33]. However, we are not aware of any broadband provider making them directly accessible to consumers.

A number of Internet services can be constructed using Diffserv. These can be matched, for example, to real-time inelastic traffic, non-realtime inelastic traffic (e.g., data backups), low-value elastic traffic, or highly critical traffic (e.g., for network control or public safety applications).

vi. Real-time, inelastic service suitable for streaming

As an example of a valuable Internet service other than Best Effort, consider a Streaming service which is designed to transport real-time inelastic video streams. This implies the following:

a. each data flow has a characteristic data rate

b. the characteristic rate forms a "traffic specification", which is specified by the source of the data flow and made known to the network

c. a data flow that transmits at a rate not greater than the traffic specification (i.e., a "conforming" data flow) should experience no packet loss due to congestion, and small, statistically bounded amounts of variable delay

d. the network meters the data flow at ingress, characterizing packets as conforming to the traffic specification, non-conforming and to be marked as such, or non-conforming and to be discarded

e. each packet is temporarily stored in a buffer at each successive hop (i.e., switch or router) along the path of the flow. Packet buffers are organized into FIFO queues. These queues are shared by multiple flows which use this Streaming service and have a common next hop. Packets within a queue are transmitted in the order in which they are received, regardless of which flow they belong to

[30] c.f., "PacketCable 2.0 Quality of Service Technical Report" http://www.cablelabs.com/wp-content/uploads/specdocs/PKT-TR-QOS-C01-070925.pdf

[31] cite – DOCSIS 3.0

[32] cite – Broadband Forum TR--121

[33] cite – Broadband Forum TR-168

f. streaming service queues are served preferentially to best effort queues

g. the depth of each queue is monitored, and one or more packets that are marked as non-conforming are discarded if the queue exceeds a pre-determined threshold,

h. the network determines a priori whether it has enough capacity to carry a new data flow; if not, the data flow is not "admitted"

i. the source's data rate is not assumed to adjust to congestion in the network

These behaviors are well described and analyzed in the networking research literature[34]. Note that (f) could be considered to be a "red flag" for prohibited discrimination in the Open Internet order.

Engineers will recognize the Streaming service as having "open loop control". By contrast, Best Effort has a nominal "closed loop control" system through TCP.

vii. Bandwidth Reservation and Admission Control are necessary for real-time inelastic services

The Streaming service described above is realized using a combination of mechanisms that are not needed for Best Effort service[35]. These stabilize network performance. One of these, "bandwidth reservation", coupled with "admission control", is particular importance to the issues raised in this proceeding.

Bandwidth reservation means that the network designates a portion of the capacity of each link along the source to destination path as being preferentially assigned to a flow. The size of that portion is a function of the traffic specification. To oversimplify, bandwidth reservation assures that packets belonging to a conforming flow can cross the network almost unimpeded by other traffic. Typically, if a conforming inelastic flow sends at a rate less than the traffic specification, the unused capacity is available for other flows that are short-term non-conforming, or for best effort flows.

No more than a link's total capacity can be reserved for inelastic flows. Therefore, the network explicitly refuses to admit new inelastic flows when that capacity is reached. Networks typically set aside a "pool" of capacity for use by best effort flows, further restricting the capacity available for

[34] the seminal work in this area is: Parekh, A.K. ; Gallager, R.G. "A generalized processor sharing approach to flow control in integrated services networks-the single node case" IEEE/ACM Transactions on Networking, Volume: 1, Issue: 3 (June, 1993).

[35] RFC 2475, RFC 4594

inelastic flows. Admission control and bandwidth reservation can be actuated administratively, or by a control protocol such as the Resource Reservation Protocol (RSVP)[36].

viii. Bandwidth reservation cannot be free

Unrestricted bandwidth reservation results in perverse incentives. Users, given the opportunity to seize capacity for their preferential use, will greedily grab as much as possible. Arbitrarily large reservations will be made without regard to whether they will actually be used, or to their consequences for other users. This is a Tragedy of the Commons.

Researchers have shown that some form of graduated pricing is necessary to maximize utility (consumer satisfaction) in multi-service networks[37].

> …monetary incentives can be used to induce users to choose the appropriate service classes and spread the resulting "utility" gains among all the users.

The Internet has long had flat-rate pricing for Best Effort service. Consumers have resisted attempts to introduce metered pricing. However, it seems likely that they might be willing to pay something for an "upgrade" to a streaming video that assures them of better video quality. This hypothesis is best tested by the market.

For the balance of this discussion, we will refer to services that reserve capacity as "Premium" services.

B. The Internet can be shared effectively between multiple services

As discussed above, the networking technology community has long understood that multiple services with different performance objectives and traffic handling paradigms can coexist on the same network[38]. This leads to the public policy question of how to ensure this is done only in a pro-competitive, pro-consumer fashion.

The policy objectives should be:

a. to allow broadband providers and ISPs to offer services that are optimized to classes of applications in ways that benefit consumers

b. to ensure that doing so does not degrade best effort service

[36] RFC 2205
[37] Cocchi et al., "A Study of Priority Pricing in Multiple Service Class Networks", ACM Sigcomm 1991
[38] see Parekh, infra.

c. to require that such services are provided in a manner which is free of discrimination, self-dealing, or unjust pricing

i. Not degrading Best Effort service

Commenters[39] have expressed concern that preferential treatment for some traffic will result in degraded performance for other traffic. Thus, they posit that service differentiation will necessarily disadvantage new applications, smaller entities or entities that lack special relationships with broadband providers. This is based on an underlying assumption that capacity allocation must be a zero-sum game. We respectfully disagree.

It is true that, for example, scheduling a packet belonging a streaming flow ahead of an earlier packet belonging to a Best Effort flow will delay the latter. It is also true that systematically doing so might cause Best Effort queues to overflow sooner than they would if all packets were treated in strict FIFO fashion. However, good traffic engineering practices can render these effects to be insignificant.

First, consider the timescales involved. Typical links between nodes on the Internet operate at nominal rates of 10 Gbit/s, with many backbone links at 40 or 100 Gbit/s. On a 10 Gbit/s link, a 1500 byte streaming packet that "cuts ahead of" a Best Effort packet will delay the latter by 1.2 microseconds (µs). Similarly, current broadband access networks have shared downstream links that operate at rates between 152 Mbit/s[40] and 2.4 Gbit/s[41]. For the former link, the additional wait is 66 µs. Such delays are trivial. Human senses are incapable of discerning them, and TCP operates over timescales which are two or more orders of magnitude greater.

Meaningful degradation only occurs if aggregate Best Effort traffic exceeds the capacity of one or more Best Effort pool along the path a flow. This degradation takes the form of delay and packet loss. Note that this is also what happens in purely best effort networks when traffic exceeds link capacity. From the perspective of Best Effort flows, it is the same as reducing the capacity of the link without increasing the time required to transmit a packet.

Two mechanisms are used to prevent Best Effort congestion in a multi-service network. First, admission control is designed to allocate reserved bandwidth from a pre-defined reserved bandwidth pool, which is disjoint from the pre-defined best effort pool. If, for example, a new Streaming flow attempts to

[39] e.g., cite Public Knowledge, Free Press
[40] for DOCSIS 3.0 downstream with four-channel bonding groups
[41] for GPON downstream; note that standards exist for 10 Gbit/s-class PONs, but are not yet widely deployed

reserve more bandwidth than is available in the reserved pool, it is blocked (i.e., not admitted).[42] As a result, best effort flows cooperate to share a capacity at least equal to the size of the best effort pool.

Second, modern routers and switches have sophisticated packet "schedulers". A scheduler is the function which determines which packet will be transmitted next on an outgoing link, based on a "service discipline" and configured policy. For example, "Weighted Fair Queuing"[43] is a commonly implemented scheduling discipline. It can be configured with "weights" for each queue, such that a Best Effort queue is assured of a particular portion of the link's capacity.

Using these two mechanisms, network engineers can manage the amount of capacity that is always available for best effort traffic. This, in combination with projected traffic statistics, can be used to assure that Best Effort flows can achieve at least a target Minimum rate.

Furthermore, segregating elastic traffic from inelastic traffic can have salubrious effect on Best Effort service performance. Recollect that the TCP protocol dynamically adjusts transmission rate, based on congestion signals detected at the receiver. By "conforming" to the rules embodied in TCP, elastic flows effectively cooperate to share network capacity in a nominally "fair" fashion. When TCP non-conforming inelastic flows are mixed with TCP conforming elastic flows, they do not cooperate. Thus, the conforming flows will adjust to divide amongst themselves what capacity remains, while non-conforming flows suffer occasional packet drops. By contrast, satisfaction is maximized if the network segregates non-TCP conforming from TCP conforming flows, sets aside a pool of capacity for the latter to share, and admits only as much of the former as it can reserve capacity for.

ii. Avoiding perverse incentives

As we have shown, broadband providers and ISPs can adjust the size of bandwidth pools for Best Effort and for other services to meet performance objectives. However, if Best Effort service is charged at a flat rate and bandwidth reservation is priced at a premium, revenue is maximized by admitting any and all premium, non-best effort flows as requested. This means starving Best Effort, which leads to the concerns expressed, e.g., by Free Press, about "toll lanes"[44]. The Commission's problem is thus to counterbalance these perverse incentives, under what statutory authority they have, preferably in the simplest way possible.

---

[42] preemption is also possible, e.g., for public safety applications

[43] Stiliadis, D.; Varma, A. "Latency-rate servers: A general model for analysis of traffic scheduling algorithms". IEEE/ACM Transactions on Networking 6 (5): 611.(October, 1998).

[44] cite

We suggest that the Commission's best leverage is through disclosure. For example, a Rule could require broadband providers who offer premium services to disclose the following engineering guidelines and measurements:

a. minimum "statistical gain" ("overbooking") factors for Best Effort traffic on shared resources at various points in the network

b. engineering targets for Best Effort packet drop rate and 95th percentile queue depth at network elements during busy hour

c. measured best effort packet drop rate and 95th percentile queue depth during busy hour

d. number of minutes per reporting period that engineering targets have been exceeded

The details of the information to be disclosed would need to be developed by experts through a Rulemaking, as we discuss further below. Measurement collection is an integral part of providers' existing equipment, and if measurements are selected carefully, they could be collected with minimum burden.

Given data such as the above, it is possible to assess whether premium services are degrading Best Effort performance. Consumers in markets with more than one broadband provider can use such information when selecting providers, incenting competition. The Commission and State regulators could use it as evidence for enforcement in event of severe and chronic degradation. Such information would also help form a record to inform future Commission proceedings. In addition, providers would have a standardized way to assess their own performance against their peers, adjusting engineering guidelines and/or making capital expenditures appropriately.

C. Ensuring that Premium services are pro-competitive

Commenters, e.g., Free Press[45] and Common Knowledge[46], express concern that "discrimination" will be used in an anti-competitive fashion. In particular, broadband providers could obtain competitive advantage by giving affiliated[47] video providers' data flows better performance than those of non-affiliated providers[48]. Alternatively, they could provide special discounts to affiliated video providers,

[45] cite
[46] cite
[47] "affiliated" is taken in the broadest sense, including entities with which the broadband provider or ISP has common ownership, strategic partnership, investment, joint marketing agreement or other special relationships
[48] cite. Note that broadband providers already afford special treatment to their own television and VoIP services.

allowing them to compete more effectively on price. We agree that broadband providers and ISPs have these perverse incentives, but also believe they can be regulated to prevent abuse.

The strongest protections would take the form of common carriage requirements. Since the Court has taken that off the table, the Commission must fall back on its authority under Section 706a.

A pro-competitive regulatory regime for these premium services must be premised on a requirement that if any such service is provided, it must be offered to all comers on a reasonable and non-discriminatory basis. It also must be provided to non-affiliated entities at the same pricing, terms, conditions and performance as it is to affiliated entities.

Again, disclosure and transparency are key. Enabling Rules should require broadband providers who offer a premium service to disclose functions and features, pricing, terms and conditions, interfaces and protocols, service level agreements, and all other relevant details. This need not take the form of a Tariff filed with the FCC; a public web page would suffice. Such disclosure would be minimally burdensome on broadband providers unless they intended to engage in abusive practices.

In addition, Rules should require disclosure of accounting details concerning transfer of premium services to affiliated entities. Enough information should to be provided such that self-dealing can be discerned from the entities' quarterly financial statements.

i. Premium services need not inhibit independent content providers or innovative applications

Commenters have also expressed concern that cost of premium services might create barriers to independent entities e.g., not-for-profit content providers, entrepreneurs, developers of novel applications or experimenters. This need not be a problem.

Under the regime we outline, best effort service, charged at flat rates, will continue to be the default service. We have shown that premium services can be offered without noticeably degrading best effort service performance. Thus, if an independent entity chose to use best effort service – regardless of whether the application were elastic or inelastic – they would not be harmed by reservation.

In addition, availability of premium services might be beneficial to independent entities. They might, for example, enable novel inelastic applications which are particularly sensitive to delay variation and packet loss. Independent entities might also offer a "free" basic version of their product or service using best effort service, and a better experience for pay using a premium service; the end user would simply have to decide whether or not to pay for better quality. This is consistent with many content and application provider business models, and provides another incentive for consumers to pay for content, applications and services that they particularly value.

ii. The Internet can support services which are optimized to classes of applications

In this discussion, we have used a premium Internet service optimized for entertainment quality video as an example. The Internet's architecture provides mechanisms to support services which are optimized to other applications, services or content[49].

For example, Best Effort service is not particularly optimal for VoIP telephony. It can be used, but user experience is improved if delay variation is bounded, and public safety requires underlying network capabilities to support E911 and Lawful Intercept. Broadband providers' own managed networks provide such a service in support of VoIP-based phone services to consumers[50]. We suggest that there might be substantial benefit to their making the same service available to competitive providers who wish to use it.

As another example, Best Effort service is not particularly optimized for sensor/actuator applications – also known as "Internet of Things" or "M2M". Such applications typically send short messages at relatively infrequent intervals and are sensitive to packet loss. The closed-loop control system inherent in TCP does not operate on traffic flows of this nature. Using the Diffserv "toolkit", it would be easy build a service optimized for such low rate, loss sensitive, inelastic traffic.

V. DISCLOSURE AND TRANSPARENCY ARE POWERFUL TOOLS FOR THE OPEN INTERNET

Justice Brandeis famously observed: "Publicity is justly commended as a remedy for social and industrial diseases. Sunlight is said to be the best of disinfectants; electric light the most efficient policeman"[51]. The Open Internet Order requires that "Fixed and mobile broadband providers must disclose the network management practices, performance characteristics, and terms and conditions of their broadband services". The Court upheld this part of the Rule, and indeed cited it as a valuable protection for consumers and competition[52]. We agree that transparency and disclosure are powerful levers for

[49] c.f., RFC 4594
[50] cite, e.g., PacketCable
[51] Louis Brandeis, *Other People's Money—and How Bankers Use It* (1914)
[52] cite Opinion, in particular footnotes 3 and 8

regulators, other public officials, "edge providers", public interest groups and individual consumers. It therefore falls to the Commission to assert this authority to enforce openness.

D. Disclosure can expose anti-competitive discrimination

Blocking or deliberately degrading performance of a specific content source or application requires an affirmative action by the broadband provider. Specifically, one or more network element must be configured to:

a. classify all data flows, based on information carried in each packet such as IP Source and Destination Address, Protocol, Differentiated Services Code Point, and Port Number,

b. identify packets belonging to those data flows which the provider wishes to discriminate against, and,

c. perform one or more action against some or all of those packets, according to some policy. For example, the network element might discard some or all the identified packets, apply a "rate shaper" or "rate conditioner", mark them as discardable in event of congestion, or enqueue them in a low priority queue.

Note that all of these steps are part of normal packet forwarding in the Internet. Some commenters seem to want to prohibit them[53] altogether. We vigorously disagree. They are indispensable tools for what the Open Internet order calls "network management". In particular, they are a first line of defense against malicious network attacks. However, like many tools, they may be applied for benign or harmful purposes.

We emphasize that any question of the form "is traffic to or from x being discriminated against" can be answered by knowing how the provider has configured packet classifiers and policies. Edge providers, consumer or others who suspect that content or applications are being unfairly discriminated against would be well armed by having access to this information. This would have a powerful deterrent effect against abusive behavior by Broadband providers. It would also protect Broadband providers against false allegations of abuse.

Therefore, we suggest that the Commission consider a Rule by which interested parties might easily obtain, on demand, details of packet classification and policies as they affect identifiable flows. Such a Rule would have to take into account concerns for subscribers' privacy and network integrity. Providers

[53] cite

would have to be required to certify that the information provided matches with actual configuration records of network elements. The details of such a Rule would need to be determined through the Commission's Rulemaking process.

E. Disclosure can address neglect of Best Effort service

As we noted above, disclosure is a remedy for Commenters' concerns that broadband providers will profit by setting up "toll roads", and highly disadvantage users of the flat-rate Best Effort service through neglect. Here, we must again digress into a technical discussion.

Statistical sharing is a fundamental property of packet switched networks like the Internet. Resources, in the form of physical and logical links and parts of network elements, are time-shared amongst numerous data flows. Sharing is flexible in nature. At each "node" (i.e., network element) packets with various sources and/or destinations are received, stored in buffers, and "forwarded", or transmitted on a selected link. The order that packets appear on an outgoing link is determined both by their arrival order and a scheduling policy.

Few sources emit packets on a periodic basis; emulated circuit switched services are the only significant exception. Instead, the intervals between packets vary arbitrarily. Furthermore, sources can be idle for significant periods of time. Packet networks are designed to take advantage of this property to efficiently utilize capacity of shared resources. Each subscriber has a notional transmit and receive rate. Network engineers typically allocate subscribers to each resource well in excess of the resource's capacity to carry all traffic at the notional rate. This practice is somewhat analogous to the practice of overbooking airline reservations. The benefit of doing so is called "statistical gain" or "statistical multiplexing gain". Statistical gain is an important element of network economics. Without it, networks would have to be massively overprovisioned, and the resulting cost reflected in service pricing.

In a well-designed network, statistical gain rarely results in noticeable performance degradation. Networks are commonly mathematically modeled as queueing systems. Inter-packet arrival times, waiting times and queue depths each vary according to a statistical distribution. Since buffers are of finite size, packets must be discarded if a queue gets too long. Packet discard events over time are measured as a "discard rate". User experience is directly related to waiting time and discard rate, which are in turn related to arrival rates. For TCP-responsive traffic, "throughput" (i.e., rate of useful data transmission) is also a function of waiting time and discard rate. Thus, these statistics are measures of network performance, and an important predictor of end-user satisfaction.

Networks are typically designed to meet specific performance objectives. Most often, this is done by rule-of-thumb engineering rules. For example, a broadband provider might decide to share a 1 Gigabit per second link amongst 200 customers, each of which has a service advertised as “up to 100 Megabits per second”. The ratio of the sum of the notional service rates to actual capacity in a resource is called “statistical gain factor”, which is 20:1 in our example. Engineers can determine the projected relationship between design performance objectives and statistical gain factor by simulation, by measurement in similar systems, or “seat of the pants”. The relationship between statistical gain factor and actual network performance is determined by measurement.

Ideally, network providers would continuously validate measured performance against performance objectives, and add resources as necessary when the former is found deficient. This might not necessarily mean immediately adding capacity every time the packet drop rate or mean delay of a resource exceeds a threshold. It does mean that chronic failure to meet performance objectives requires new capacity.

With that explanation in mind, we suggest that the Commission consider requiring broadband providers to disclose the following:

a. target performance objectives at each network element, especially for Best Effort service,

b. engineering design rules for network facilities, and

c. measured performance at each network element (or deviation from performance objectives)

Such disclosure would be somewhat analogous to an employee performance review. It would be enormously useful to regulators, local governments, consumer advocates and individual consumers for determining how well (or poorly) a broadband provider is performing. It would quantitatively expose poor service and bring competitive or regulatory pressure to bear on low performing broadband providers to invest in necessary upgrades. It would also be helpful to broadband providers in defending themselves against unfair complaints.

The details of such a Rule needs the attention of subject matter experts, as input to the Rulemaking process. In particular, work would be needed to select meaningful performance metrics that can be measured using capabilities in existing network elements, and to determine the granularity and form of disclosure.

VI. CONCLUSION

The Internet's architecture was designed in the late 1970s and early 1980s as the ARPAnet project, a research network for use by researchers. It has evolved through continuous tinkering, under the aegis of open processes. The Internet is called upon to do many things it was not originally designed for. In particular, Video over IP, which was at best a theoretical possibility at the time the Internet was architected, now comprises most of the Internet's traffic[54]. Thus, we do not agree with commenters who wish to freeze a platonic ideal of a flat-rate best effort Internet in amber, never to be augmented or improved.

At the same time, we do not agree with commenters who seek to have the Commission abandon all but *de minimus* oversight. The research community that created and used the ARPAnet was self-policing and cooperative. The Internet's original architecture was not designed to provide commercial service; indeed, pecuniary interest was actively discouraged. Presently, what competition exists in broadband access markets is inadequate to impose market discipline. Such discipline is necessary not only to prevent consumer harms, but also to ensure that other entities in the Internet's ecosystem are not harmed. Thus, the Commission must actively regulate those market segments of the Internet that lack effective competition.

We have proposed what we believe to be a pragmatic middle way. In summary, the Commission should:

a. eliminate distractions

b. divide and conquer

c. concretely and precisely define issues

d. move swiftly to close non-contentious issues

e. require more disclosure, particularly of engineering guidelines, performance objectives, actual performance, pricing, and intra-company transfers

f. encourage broadband providers to offer services in addition to the default Best Effort service, conditional on principles of fair, reasonable and non-discriminatory availability and pricing, and measurably negligible degradation of Best Effort service

---

[54] cite

While this is not exactly what all the parties might prefer, it represents the basis of a good compromise, and we believe it serves consumers.