Before the

Federal Communications Commission

Washington, D.C. 20554

In the Matter of the Open Internet Order Remand

GN Docket No. 14-28

July 16, 2014

Comments of NetAccess Futures

NetAccess Futures, through its Principal, Daniel B. Grossman, hereby submits these comments in response to the NPRM adopted May 15, 2014, in the above captioned matter.

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 instilled respect for their role in maintaining competitive markets.

We have not engaged counsel to prepare these *pro bono* comments, which therefore no doubt deviate from normative form. We have also not footnoted as extensively as we might, and most footnotes are incomplete. Further, we have not developed all of our arguments as fully as we might have preferred. Nonetheless, 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

NetAccess Futures previously filed comments in this matter[1]. We advocated a pragmatic middle way. We encouraged the Commission to maintain focus in this proceeding on the core issues raised by *Verizon vs FCC*, while addressing related matters separately. We further noted that the non-blocking principle of the Open Internet order, while contested by Verizon, is not in itself controversial. We explained the usefulness of services based on differentiated treatment of traffic. We then observed that the policy risks of such services – degradation of the normative "Best Effort" service, and certain anti-competitive behaviors – can be mitigated through appropriate regulation.

We are gratified to note that the Chairman, in this NPRM and in his public statements, has adopted many of the ideas proposed in our previous comments. In particular, the Commission has begun to separately address peering and interconnection issues which have been conflated with Open Internet concerns[2]. Further, the NPRM presents a framework for differentiated traffic handling consistent with our proposal.

By these comments, we hope to assist the Commission by reframing a toxic debate.

II. THE COMMISSION'S SITUATION

The Commission is presented with a fine mess. Several public interest parties[3] have stirred a firestorm of public opinion and comment. The Commission rightly acknowledges public interest in this matter. Members of Congress have weighed in on both sides[4]. The Commission must respond to the will of Congress, but deep partisan and ideological divisions leave Congressional intent unclear. The Commission must work within the authority granted by the Telecommunications Act of 1996. Yet the Act, and its basis in the Commission's *Computer II* order, has become obsolete. The Commission must therefore fit regulation of broadband Internet services to a Bed of Procrustes.

The Commission must somehow balance the interests of all parties, within the possibilities and constraints of the Internet technology and economy. We emphasize that these Proceedings are not a referendum. Public opinion, while certainly intense, is not well informed[5]. We are confident that the Commission will

[1] Comments of NetAccess Futures, GN 14-28, April 3, 2014.
[2] Public notice TBA; Blog post TBA
[3] Save the Internet, Free Press, Common Knowledge, ACLU, Schenick etc. TBA
[4] c.f. HR.4070, S.1981, House Communications and Technology Subcommittee Oversight Hearing, May 20,2014
[5] See blog posting, "Democracy is Messy", http://www.netaccessfutures.net/network-neutrality/democracy-is-messy/

base new Rules on objective evidence. However, we are concerned that thoughtful, and well-argued comments will be shouted down by the mob. Indeed, we fear that our own comments might be lost in the noise.

Too much public comment in this has been exceptionally vituperative. We take exception to ad-hominem attacks on the character and motives of the Chairman (in particular), the other Commissioners and staff. We are confident that they follow their ethical obligations to act in their best judgment on behalf of their clients - the American people.

A. The record needs to be corrected

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.

B. Discourse Suffers From an Inapt Metaphor

The public misunderstands the technical and business issues at stake in this proceeding. In particular, they have been told of a "fast lane", or sometimes a "toll lane"[6] in the Internet. Much discourse thus turns on a metaphor that is flat-out wrong.

The "Information Superhighway" metaphor is credited to former Senator Al Gore Jr[7]. It was a tribute to his father, Sen. Al Gore Sr., who was instrumental in the legislation that created the Interstate Highway system[8]. It was also an astonishingly prescient prediction of the impact which the Internet would have on our daily lives. Mr. Gore is no technologist. He needed some way to express the notion of a ubiquitous,

[6] c.f. "Save the Internet"
[7] cite
[8] cite

richly interconnected data network as critical infrastructure for the 21st Century. The highway metaphor served his purpose. We doubt he intended it to be taken too literally.

The problem is that the Internet not like a like a highway. The behavior of a road system in carrying individual cars is completely unlike the behavior of the Internet, which carries a duality of individual packets and flows. Modern Physics teaches that light simultaneously has a wave nature and a particle nature; similarly, traffic on the Internet simultaneously has a packet nature and a flow nature.

As an example, if a highway becomes congested, all cars slow down or stop, cars back up, and the resulting traffic jam grows indefinitely until the congestion clears. The Internet handles congestion by dropping packets, with the expectation that the receiver will detect missing packets from each flow, and take their absence as an indication that a congestion event has occurred. The receiver is then expected to instruct the sender to send fewer packets belonging to the flow at a time. Now, imagine a highway that handled congestion this way. Would it have artillery pieces at intersections to blow up random cars? Unless cars traveled in something analogous to a flow, how would a destination know that a car had gone missing, or signal back to an origin that it should dispatch fewer cars at a time?

And that's just one of the Internet's behaviors. If, to extend this thought experiment, one were to imagine a transportation system that behaved like the Internet, it would be truly bizarre.

Other metaphors fail as well. The late Sen. Ted Stevens was roundly ridiculed for comparing the Internet to "a series of tubes"[9]. That metaphor holds no better – but no worse – than the highway metaphor.

We emphasize that the Internet behaves like nothing in everyday experience. Thus, reasoning by the highway metaphor has been a cardinal fallacy in this debate. If public interest and other parties started to think about the Internet on its own terms, perhaps they would not find the direction of the NPRM to be as threatening.

C. "Paid prioritization" reflects deep misunderstanding

The Commission needs to reframe the debate.

The vernacular notion of a "fast lane" apparently is formally understood as "paid prioritization"[10] . This is an inaccurate and prejudicial characterization. "Paid prioritization" is fully described the following network behavior:

[9] cite
[10] cite

> Edge Provider A contracts with Broadband Provider B for strict priority treatment of certain data flows, e.g., from A to customer C. Edge provider D has no such arrangement. When a packet from A to C arrives at a router in B's network, it is scheduled for transmission ahead of all other packets, e.g., a previously arrived packet from D to customer E. It cuts to the head of the line. A may send packets to C at an arbitrary rate. Over time, the rate of flow from D to E is reduced in order to satisfy the flow from A to C. If <A-C> and <D-E> are the only two flows traversing a congested resource having rate R, the throughput of <D-E> = R - <A-C>. Thus, A obtains superior service at the expense of C.

This is not a particularly useful behavior. If Edge Provider D contracts with Broadband Provider B for strict priority treatment, <A-C> and <D-E> have both paid for the same service they would have received *ex ante*. Strict priority does not assure that A's customers receive service at any particular rate, delay bound or loss rate. Thus, the contract between A and B cannot have meaningful service level agreements (SLAs), nor can A verify that they are receiving the contracted service. Further, utility is not maximized by such a policy: the service perceived by customer C is subject to diminishing returns on the increased rate, while customer D is needlessly frustrated.

As we discussed in our previous comments[11], and further discuss below, traffic belonging to some applications is best treated according to a policy other than "Best Effort". The Internet's architecture includes a rich "tool kit" of mechanisms for realizing differentiated for differentiated packet treatment[12]. Some of these provide queueing and scheduling disciplines other than the assumed strict First-in, First-out model implicitly assumed by Best Effort. Strict priority is one of those mechanisms; however the IETF retained it mainly for backward compatibility with implementations of IP prior to 1999.[13]

Also as we discussed in our previous comments[14] and below, another part of the "tool kit" is the resource reservation mechanism. Performance SLAs are realized, in part, by ensuring that sufficient capacity is available at each resource along the path of the assured flow. That capacity is reserved for the flow's preferential, but non-exclusive, use. All flows are incented to reserve as much capacity for themselves as possible, leading to a Tragedy of the Commons. Thus, reservation must be priced (i.e., by the Broadband

[11] cite
[12] RFC 2475
[13] RFC 2474
[14]

Provider) to be commensurate with its economic utility. This incents customers to reserve no more bandwidth than they need[15].

Therefore, the proper question is not "Should paid priority be prohibited?" Instead, it should be "Should network services other than Best Effort be prohibited? Should charging models other than flat rate be prohibited?"

We urge the Commission to eliminate the use of the term "Paid Priority" in the proposed rule, and limit its use in future Orders.

D. The notion of "services" clarifies the issues at hand and reframes the debate

The IETF defines a "Service" as:

> A description of the overall treatment of (a subset of) a customer's traffic across a particular domain, across a set of interconnected… domains, or end-to-end….[16]

Best Effort is an example of a service. Its treatment of customer traffic is well understood. We discuss its characteristics, in depth, below.

Best Effort Service is not the exclusive service of the Internet. Other services can be created that treat packets in alternative ways. This is the heart of the debate: are other services allowed or not.

This is more than mere technical pedantry. The notion of a default Best Effort service, which can be supplemented by specialized services, draws directly on the Internet's architectural framework. It broadens the discussion from one particular specialized service – Absolute priority – to all services which could be offered by a broadband provider. Absent facts to the contrary, we would likely consider an Absolute priority service to be "commercially unreasonable"; however, there are potential special services which are commercially reasonable. The debate suffers from ambiguity between Absolute priority service in particular, and special services in general, and from narratives which conflate the former with the latter.

For the balance of these comments, we the follow the lead of the Commission[17] in using the term "Specialized Services" to means services (under the IETF definition) other than Best Effort. The proposed

[15] Cocchi, Ron, et al. "A study of priority pricing in multiple service class networks." ACM SIGCOMM Computer Communication Review. Vol. 21. No. 4. ACM, 1991.
[16] RFC 2474 at 2.
[17] NPRM at 60; Open Internet Order, 25 FCC Rcd at 17909, 17965-66, paras. 7, 112-1

Rule[18] uses the term "Priority service". This, in particular, should be replaced by "Special service" or some other term.

The Commission should reframe the proposed Rule to draw on the notion of "service", particularly of Best Effort service and Specialized services.

III. OUTLINE OF OUR POSITION

We intend to respond to as many of the Commission's requests for comments as our time and expertise allow. However, our responses are best understood in the context of our overall framework.

1. We mostly concur with the Chairman's proposal, as reflected in the NPRM.
2. We take the societal value of the Internet as a given that requires no elaboration. The Internet's facilities are critical infrastructure and a utility (with a small 'u').
3. There is insufficient competition in local broadband access markets to assert market discipline. Broadband providers have (but do not necessarily act upon) perverse incentives, supported by lack of market discipline. Specifically, they are in a position to engage in anti-competitive and anti-consumer practices, extraction of monopoly rents, self-dealing and (arguably) suppression of public discourse. The Commission's Rules and enforcement process must therefore substitute for market forces in prohibiting such abuses.
4. "The Internet" represents a value chain of consumers, network providers, edge providers, content providers and others. Allocation of consumer revenue among the rest of the value chain is a zero-sum game. The Commission will frequently be called upon to adjudicate tussles over cost and revenue[19].
5. The Internet is a Network of Networks. While nobody "owns the Internet", some entity owns each of its constituent networks. Those entities are often for-profit corporations. Capital and operating expenses for building, operating, maintaining and upgrading networks – particularly raccess network infrastructure – are immense. The Internet's business model must support adequate return-on-investment to incent these investments.
6. The Internet's architecture is not engraved on stone tablets, according to some platonic ideal, forever immutable. It evolves in response to identified engineering and business problems. The way the Internet actually works is very different from the way it is generally perceived to work.

[18] NPRM at Appendix A, §8.3
[19] c.f., Netflix comments

7. Point-of-view based blocking and degradation of traffic is odious and unacceptable. The Commission must be vigilant against anti-competitive blocking and degradation.
8. The Internet is a content agnostic platform for all forms of communications. However, the technical and economic characteristics of various communications varies widely. Traffic handling behaviors are optimal for some applications, content and protocols, and are sub-optimal (or even inadequate) for others. The Internet can and should handle traffic in the most optimal fashion.
9. The Internet's canonical and default traffic handling behavior is "Best Effort", with flat rate charging.
10. Other traffic handling behaviors and charging models can coexist with Best Effort/flat rate. With proper engineering, coexistence does not meaningfully degrading Best Effort service.
11. Each network element in the Internet is subject to congestion when offered load exceeds capacity. Congestion manifests itself as delay and packet loss. The Internet's canonical congestion control system attempts to mitigate congestion by voluntary, cooperative sharing of bottleneck capacity.
12. Best Effort service must treat traffic from all sources to all destinations roughly "fairly", without commercial discrimination. Excess delay and packet dropping should occur only during periods of congestion.
13. The Internet's dynamics are complex, incompletely understood, and often counter-intuitive. Differentiated traffic handling can, in some circumstances, improve overall performance.
14. Broadband providers should be permitted to offer "Specialized Services" in addition to Best Effort, subject to competitive and consumer safeguards.
15. The above safeguards must address two problems: resource starvation of Best Effort service; and unreasonable discrimination against disfavored edge providers or unaffiliated entities.
16. Resource starvation can best be addressed by disclosure and enforcement. Broadband providers should be required to disclose design rules and engineering objectives for Best Effort service and any deviations from those objectives. Performance measurements used by broadband providers should be made available to consumers, edge providers, networking professionals and regulators, subject to safeguards for privacy and network integrity. Admission control and bandwidth reservation policies should be disclosed. Traffic classification terms and actions should be disclosed, subject to safeguards for privacy, lawful intercept, national security and network integrity.
17. Pricing, terms and conditions for any Special Services offered by a broadband provider should be disclosed. Refusal to offer a Special Service to a similarly situated entity, under substantially equal

pricing, terms and conditions, should be considered a commercially unreasonable anti-competitive practice.

18. The Commission should be vigilant against egregious practices and prepared to take enforcement action when needed.

This proceeding should remain limited in scope to treatment of packets and data flows in the Internet. Other considerations, especially interconnection and inter-provider compensation, should be subject of other proceedings.

IV. POLICY GOALS

The Commission's fundamental question is: "What is the right public policy to ensure that the Internet remains open?"[20] We suggest the following high-level policy goals.

1. Maximization of welfare among players in the Internet's ecosystem, especially consumers.
2. Reinforcing the recognized benefits of the Internet's present service and business model as a general communications platform for commerce, entertainment, automation, private dialog, public discourse and civic engagement.
3. Encouraging the Internet's service and business model to be extended, to enable new applications, optimize performance and maintain orderly growth, in a pro-competitive, pro-consumer fashion
4. Eliminating perverse anti-competitive and anti-consumer incentives or prohibiting operators from acting upon them

Any new Open Internet order should be evaluated as to whether or not it advances these objectives.

V. THE INTERNET PLATFORM IS A *TABULA RASA*

The Commission asks for comments regarding the benefits of the Internet as platform for innovation, economic development and civic engagement. The question is anodyne; there is no doubt the Internet has had tremendous impact on society and the economy. The proper question should be "how", not "if".

The Internet's underlying technology, business model and governance provide a lowest common denominator for ubiquitous multimedia communications. While the technology is by no means optimal

[20] NPRM at p.3

for any application, it is sufficient to enable most applications, subject to available capacity. Other technologies and organizational structures might have served this role[21]; the market selected the Internet as we know it because it was the path of least resistance.

All societal benefits (and liabilities) of the Internet stem from nearly ubiquitous availability, interoperability, the roughly content-agnostic nature of the service, and (arguably) flat rate charging. Implementers use relatively simple applications program interfaces (APIs) to access the Internet's relatively simple Service. Applications can be developed without networking expertise; with the emergence of cloud-based services, they can also scale without networking expertise.

The Internet's basic service is "Best Effort"[22]: it tries to deliver as many packets as it can, as fast as it can, subject to resource constraints, without any guarantees. The Best Effort service, in conjunction with the congestion control functions in the Transmission Control Protocol (TCP), attempts to provide each data flow with a nominally fair share of the capacity of the bottleneck link of the path. Best Effort is a central concept for understanding the balance of these comments. While Best Effort service is sub-optimal for many applications and scenarios, it is "good enough" for most.

The Internet has thus far largely eschewed usage-based accounting and charging. End customers are charged at a flat rate, which typically varies with the advertised data rate at the network interface. This is advantageous to consumers, who can budget for their monthly broadband bills without surprises. The lack of accounting, rating and extra billing machinery reduces complexity for Internet Service Providers. The problem with flat rate charging is that the Best Effort service makes each shared resource a Commons, with no disincentive to excessive usage; thus, a Tragedy of the Commons is avoided only by capital spending. "Bandwidth Hogs" have been a problem for broadband providers, as they crowd out lower volume users and drive shorter upgrade cycles. Some operators have imposed volume caps in an effort to disincent excessive usage and allocate capital costs to those who drive them.

It has been argued[23] that a causal relationship exists between the Internet's Best Effort, flat-rate service and its success as a platform. We agree, but with the caveat that Best Effort/Flat Rate is not mutually exclusive with other service and business models

[21] cites to OSI, ISDN, and ATM
[22] cite
[23] cite

The Internet's "openness" is more precisely described as architectural decoupling between networking functions and application functions. The OSI Reference Model[24] attempted to rigorously organize data communications into "layers" of cohesive functions. Layers are loosely coupled together through simple, abstract and formally defined interfaces. Each layer offers a defined Service to the next higher layer, and uses a defined Service of the next lower layer. A lower layer neither knows nor cares about the "User Data" presented to it by any higher layer. The Internet model adopts the layering principle, although not the rigor or specific layers of OSI. The "End-to-End Principle"[25], a core tenet of the Internet, is another way to express the notion of decoupling between networking and application functions.

Decoupling does not preclude a lower layer service from offering specific, optional Quality of Service (QOS) objectives to the next higher layer. OSI explicitly recognized this[26]. The Internet's layering is less explicit about this, but as we noted[27], the Internet's protocols have always provided mechanisms that allow data to request and obtain differentiated performance. Such mechanisms do not, in themselves, compromise openness. Indeed, the NPRM notes[28] that so-called "specialized services", e.g., facilities-based VOIP, are not subject to the Open Internet order. These specialized services use differentiated performance mechanisms.

We find the public alarm about supposed threats to the Open Internet to be largely unfounded. Nothing proposed in the NPRM even remotely suggests a threat to innovation, public discourse and civic engagement. The parade of the horribles raised by public interest parties[29] misunderstands the technology and economics of the Internet, and assumes deliberate malice on the part broadband providers.

The Internet has unquestionably provided a platform for innovation in devices, applications, discovery, learning, community, private communication, commerce, content marketing and distribution. The Internet's platform is also subject to innovation, not just in "feeds and speeds", but also in optimization, performance, and business models. Some innovators in the former space[30] have adopted an attitude toward broadband providers of "innovation for me, but not for thee". While some of the latter innovations have potential to create friction with the former, a blanket prohibition of large classes

[24] cite ISO 7498
[25] Saltzer, Reed, Clark "The End-to-End Principle", cite
[26] OSI RM, cite section
[27] ref NetAccess Futures Comments
[28] NPRM at 60
[29] cite
[30] cite Internet Association

innovation in network platforms is not sound policy. The Commission's enforcement process should be adequate to address any such tussles.

VI. THE INTERNET'S ARCHITECTURE INHERENTLY PROVIDES TECHNICAL ABILITY – BUT NOT NECESSARILY INCENTIVES - TO VIOLATE OPENNESS

The Commission seeks comment on whether broadband providers have incentive and capability to violate the principles of the Open Internet order[31]. We take no position as to their incentives. We do seek to clarify the record as to the technical mechanisms that might be abused for this purpose.[32]

A. Packet Forwarding

The Internet's architecture requires that each router along the path from source to destination to undertake a prescribed series of steps[33], some of them mandatory and some optional, to forward packets. Of particular interest are:

1. Classification of individual packets by fields in the packet headers, including IP Source and Destination Address, Protocol, Differentiated Services Code Point, and Port Number. Classifiers match to patterns of Boolean strings within these fields. Further classification using application data, or deep packet inspection (DPI) is also used. A classifier database contains a set of patterns, and for each pattern, an outgoing link and a set of one or more action to be taken by the router upon packets that match the pattern.
2. Identification of flows of packets, based upon results of classification. A flow is a series of packets from a source to a destination which are inferred to be related to a single or aggregated instance of communication.
3. Traffic conditioning of flows (or groups of flows). Traffic conditioning can take the form of rate shaping (i.e., retaining packets when the flow exceeds a configured rate, and metering them out according to that rate), rate policing (i.e., discarding packets when a flow exceeds a configured rate) or metering (i.e., passively measuring the rate of each flow). Rate shaping devolves to rate policing when the number of retained packets reaches a configured limit.
4. Queueing of packets based on classification and/or flow identity. As they are received (i.e. before classification), packets are stored in data structures called "buffers". Buffers are further organized

[31] NPRM at 39
[32] See also our comments at x
[33] RFC 1812, RFC 2794

into first-in, first-out (FIFO) queues. A router typically has several queues at the input and/or output of each interface. Each buffer is added to a queue, typically based upon packet classification results.

5. Packet discard. Routers (and their component modules) can buffer a finite number of packets. When packet arrivals exceed packet departures the router must discard some packets; i.e., by returning packet buffers to the free buffer pool. Discard is normal, even desirable, behavior[34]. In particular, it is used by the TCP protocol as a congestion signal for end-to-end rate control. For this reason, current best practice is to discard packets pro-actively[35], in advance of severe congestion. Strategies for doing so are complex, and small changes to discard algorithms or their parameters can have a significant effect on congestion.
6. Packet scheduling. When an outgoing link becomes idle, the router must select which queue to service next. Service consists of removing a packet buffer from the selected queue, serializing it, and transmitting the resulting bit stream over the link. Scheduling disciplines are typically sophisticated[36], and are configured by the broadband provider in order to effectuate various objectives. In addition, a scheduler can also serve as a traffic shaper.
7. Packet marking. Internet Protocol (IP) [37] packet headers contain a field called the "Differentiated Services Code Point" (DSCP)[38], and another field called "Explicit Congestion Notification"(ECN)[39]. The former is used to alter the match and action of packet classifiers. The latter is used as a secondary network congestion signal, to be used by TCP's congestion management mechanisms. A router may alter these fields, based on packet classification, queue occupancy and/or traffic conditioning. The altered DSCP field is seen by classifiers in subsequent routers along the path. Thus, when a router alters the DSCP field, it changes forwarding actions (e.g., different queueing or discard policies) along the remainder of the path. When a router alters the ECN field, the rate of the flow is expected to be reduced

All of these mechanisms are indispensable for network function or reasonable network management.

---

[34] ref Jim Gettys
[35] RFC 7141
[36] 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).
[37] Differences between IP Version 4 and IP Version 6 are not relevant to this discussion
[38] RFC 2474, RFC 2475
[39] RFC 3168

- Classification and sometimes flow identification are needed in order to determine which actions must be taken on a packet, and to determine the "next hop" link over which the packet is to be forwarded.
- Traffic conditioning is used in conjunction with queueing and scheduling in order to effect bandwidth allocation (as we describe below). It is also used to enforce rate caps, for congestion management[40] and defense against certain kinds of attacks.
- Queuing provides for in-order transmission of packets, for fairness between flows, and for management of bandwidth.
- Discard protects the router's buffer resources against overload, and is the primary congestion signal used by TCP to adjust its rate.
- Marking simplifies packet classification at subsequent routers, and effects traffic management for particular flows.

All of these mechanisms can also be abused, to the detriment of Open Internet principles.

- Packet classification and sometimes flow identification are necessary precursors to discrimination against (or in favor of) certain packets or flows.
- Traffic conditioning can be abused, possibly in conjunction with queueing and scheduling, to throttle disfavored flows.
- Special advantaged or disadvantaged queues can be configured for discriminatory treatment.
- Packet discard can be used unfairly to favor or disfavor specific flows through TCP's congestion avoidance mechanism.
- DSCP marking can be used to indicate favored or disfavored flows, and ECN marking can be used to effect TCP congestion avoidance for disfavored flows.

Note the chain of steps in packet forwarding: classification → flow identification (optionally) → action(s). We will revisit this later.

B. Bandwidth Reservation and Admission Control

The Best Effort service is based on cooperative sharing of resources. Packet flows start and end without explicit notification to, or permission from, the network. They are assumed to adjust their rates according to TCP's slow start and congestion avoidance mechanisms, based on presence or absence of congestion signals. This behavior gives each flow an approximately fair share of the capacity of the bottleneck link.

[40] Gettys, infa.

Rate adjustment in response to congestion signaling is a form of closed loop control. No capacity is reserved for a Best Effort flow; it takes what it gets.

Services that rely, wholly or partly, upon open loop control have different mechanisms for resource sharing. Capacity (or "bandwidth") is reserved. Reservation can be dynamic, using a signaling protocol such as RSVP[41] to negotiate bandwidth at the start of a flow, and release it at the end. It can also be static. Reservation typically means preferential, rather than exclusive, access to capacity. If a flow is not using all of its reserved bandwidth, the unused capacity is available for other flows. Open loop controlled flows are defined by a traffic specification (TSpec) which characterizes the flow's rate. A traffic conditioner enforces compliance with the TSpec, delaying, dropping or marking packets in excess of the agreed rate. The reserved bandwidth is a function of the TSpec. Some services allow a flow to have a reserved component and a Best Effort component.

The sum of reservations at each resource must not exceed the capacity of that resource. The admission control function performs bandwidth reservation[42], only if sufficient capacity is available; otherwise the network must deny the flow. The flow might proceed as Best Effort, or it might be blocked.

These flows coexist with Best Effort flows. Typically, a network engineer will allocate a pool of capacity to be shared by all Best Effort flows. The capacity in this pool is not available for reservations. Best effort flows share that pool, along with any capacity that has been reserved but not in use. The capacity of the Best Effort pool must be sufficient to meet acceptable performance targets for Best Effort traffic.

As we describe below, there are numerous valid reasons for open loop controls in addition to closed loop controls. Thus, bandwidth reservation and admission control are needed in order to offer services other than Best Effort. They also are subject to anti-competitive and anti-consumer abuse by broadband providers. If flat-rate charging applies for Best Effort service, and additional pricing applies for other services, broadband providers are incented to minimize the Best Effort pool. This has the effect of starving Best Effort flows. This is the technical substance of the feared "fast lanes". Admission control is subject to commercially unreasonable discrimination, such as giving precedence to admission of favored flows (e.g., from an affiliated entity) or denying access to non-favored flows.

C. Domain Name Service

[41] cite RFC
[42] Bandwidth broker paper; RFC 4125-4128, RFC 4804

The Internet's Domain Name Service (DNS) resolves human-usable domain names (e.g., "fcc.gov") into IP addresses (e.g., 192.52.94.5). It allows applications to use human-friendly, topology independent names, while the network infrastructure operates on machine-friendly, topology dependent addresses. It also allows content to be location independent.

DNS is a distributed database system, comprised of "name servers". It is the only part of the Internet's architecture that is hierarchical, rather than peer-to-peer in structure. The database need not be consistent: a name frequently resolves to different addresses, depending on which name server is queried.

DNS is an important service provided by ISPs; in light of *Brand X*, this includes broadband providers. Typically, when a consumer device connects to its local network, it uses the Dynamic Host Configuration Protocol (DHCP) to obtain the IP address of a local primary and secondary Name Server. The source of this information is the broadband provider's DHCP server (or its proxy), which provides the addresses of the broadband provider' DNS server (or a proxy). As a result, unless the user configures their device to use a different DNS server, name resolution is controlled by the broadband service provider.

The flexibility of DNS is extremely valuable to the Internet's operation. In particular, content distribution networks (CDNs) use DNS to resolve the name of a content resource to the most optimal server that has that content. However, this flexibility can also be abused by ISPs to redirect user queries to their own services and content. Further, the largest broadband provider ISPs[43] have recently entered the CDN market, and thus compete with independent CDN providers. This might be an incentive to improperly modify databases in their DNS servers.

D. Routing

The Internet's routing system distributes topology information, which is used by each router to create a "forwarding database". The forwarding database maps destination IP addresses to a "next hop" or outgoing link to be used by packets which carry that address.

The routing system can be abused in order to send packets to favored destinations along specially engineered paths not available to disfavored packets. It can also be abused to send packets to disfavored destinations over unnecessarily long or frequently congested paths.

[43] Comcast cite, Verizon cite

VII. "SPECIALIZED SERVICES" SHOULD BE PERMITTED, SUBJECT TO PRO-COMPETITIVE CONDITIONS

The Commission asks for comments on "specialized services"[44]. We contend that this goes to the heart of the debate over so-called "Fast Lanes".

As we detail at length in our previous comments[45], there are valid reasons for differentiated treatment of different packets and flows. Other commenters[46] agree. As we further observed, the technical capabilities for doing so are widely deployed and are used in enterprises, and by broadband providers to support their own managed services.

We urge the Commission to permit broadband providers to offer specialized services, subject to competitive and consumer safeguards as we will detail below.

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

The Commission seeks comment on the Transparency rule[47]. We discussed this in depth in our previous comments[48], even using the same quotation of Justice Brandeis. We reiterate much of that discussion here.

IX. DISCLOSURE CAN EXPOSE ANTI-COMPETITIVE DISCRIMINATION

Blocking or deliberately degrading performance of a specific content source or application requires:

1. Packet classifiers and possibly flow identifiers must be configured to identify those flows which the provider wishes to discriminate against or in favor of.
2. Upon identification, action(s) configured for disfavored or favored packets or flows must be taken in the form of discriminatory queuing and scheduling, discriminatory rate shaping or policing, discriminatory discard policy, or prejudicial marking.

[44] NPRM at 60
[45] NetAccess Futures Comments at
[46] e.g., Rob Freiden, "Net Bias and the Treatment of "Mission Critical" Bits", presentation at the Future of Broadband Regulation conference, May 30, 2014
[47] NPRM at 63
[48] NetAccess Futures Comments at

We emphasize that the broadband provider must take affirmative steps to configure classification matching and consequential action in network equipment. Thus, any question of whether the provider is unfairly discriminating can be answered by knowing how the provider has configured packet classifiers and policies. Edge providers, consumers, regulators 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 defend broadband providers against unfounded 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 must take into account concerns for subscribers' privacy, network integrity, national security and lawful intercept. Providers should further be required to certify that the information provided matches with actual configuration records of network elements. Since all relevant information exists in configuration records, and such reporting can reasonably be automated, we do not anticipate that this requirement will be an unreasonable burden. We have not considered how this information can be digested into consumer-friendly form, noting the complexities involved.

A. More disclosure is better than less

The Commission asks for comment as to what information should be provided to consumers, and in what form. We note that broadband access and Internet technology is esoteric, and nearly incomprehensible to lay people. Reducing these concepts to a form that can be understood by most people is a formidable challenge. Yet people are especially frustrated by things which they do not understand. Most consumer complaints, such as those cited in the NPRM[49] , stem from lack of deep understanding. For example, many end-to-end performance problems are not under control of broadband providers. Consumers lack knowledge and tools to determine responsibility, and thus always blame their service provider. We also observe from comments[50] that consumers do not understand the service that they are they are paying for or the delineation of responsibility between their broadband provider and other parts of the Internet's ecosystem. They perhaps have signed a contract, which may describe this with incomprehensible wording, in fine print. What is needed is plain language that explains exactly what is covered by their monthly service, including elements that the service provider offers on a best-effort basis. For that reason, we commend the efforts of the OAIC.

[49] NPRM at 69, footnotes
[50] cite

At the same time, sophisticated consumers (or consumers' agents) benefit from detailed disclosure, which they may use to quickly resolved disputes with the provider or, in the worst case, to document complaints. We urge that consumers have access to the same information as edge providers.

Performance issues with Best Effort service are difficult to diagnose[51]. Indeed, as shown in the illustration in the NPRM, congestion bottlenecks in the home network, interconnection points and elsewhere are often misattributed to broadband providers. "Black Box" measurement tools, like SpeedTest[52] and traceroute[53], can provide insight to network engineers. However, greater understanding can be obtained from information collected by network elements. Such information should be made available upon request, subject to safeguards for network integrity and consumer privacy. We believe the benefits of making it available outweigh any proprietary considerations.

B. Disclosure can address degradation of Best Effort service

Disclosure is a remedy for 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[54].

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.

Packet networks are designed to take advantage of idle periods to efficiently utilize capacity of shared resources. Each subscriber has an advertised (or headline) 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 advertised rate. 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 (i.e., excessive queueing delays and packet loss rate). 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

[51] cites including Gettys
[52] cite
[53] cite
[54] see Comments of NetAccess Futures at

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 networks, 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:

1. engineering design rules for network facilities
2. target performance objectives at each network element, especially for Best Effort service
3. measured deviation from performance objectives at each network element

The information in this disclosure exists, but is presently treated as proprietary business information by broadband providers.

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. When specialized services are offered, it would expose unreasonable discrimination against Best Effort service. It would also be helpful to broadband providers in defending themselves against unjustified complaints.

X. DEFINING A MINIMUM LEVEL OF SERVICE UNDER A NO BLOCKING RULE

The Commission asks for comments[55] on several issues related to how the Rules might specify a minimum level of service under a revised No Blocking rule. Perverse incentives to degrade the flat rate Best Effort

[55] NPRM at 97

service is the most serious anti-competitive and anti-consumer risk raised by the Commission's approach. The Court's condition that any such Rule be constructed to not impose common carrier status on broadband providers[56] makes the Commission walk a very fine line. We generally agree that the approach in the NPRM[57] is on the right track.

The Court's logic on providing to the edge provider "minimum service to satisfy the rules"[58] is backwards. The obligation should be to provide at least a minimum service to consumers, not to edge providers. It is satisfied if the broadband provider forwards packets and flows from its point-of-interconnection to the network interface (and vice-versa) with acceptable performance. Congestion at points of interconnection between broadband providers and edge providers, CDNs or transit networks should be outside the scope of this rule, to be addressed in a separate Proceeding.

The Commission further asks for comments on how to define, monitor and enforce a minimum level of access, and presents three possible approaches, which we discuss below.

XI. BEST EFFORT SERVICE

"Best Effort" is default service of the Internet[59]. It implies the following:

- The source of a data flow may transmit packets freely, without any enforced time or rate constraint
- 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. Packets waiting on a queue are delayed until previous packets have been forwarded.
- The depth of each queue is monitored, and if it exceeds some threshold, packets are discarded according to some policy
- 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

[56] *Verizon* at 658
[57] NPRM at 98
[58] NPRM at 99
[59] RFC 2474

- If every data flow responds to congestion signals by adjusting its rate as required by TCP[60], the rates of all flows will converge to an approximate "fair" share of each link along their respective paths
- 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
- Flat rate charging applies, possibly with limits (e.g., volume caps).

This is understood by the network technology community as the default and canonical service of the Internet. It the behavior that consumers, edge providers, developers and others expect. It is the sole behavior of the Open Internet as understood by many commenters. Therefore, the Commission should regard as presumptively commercially unreasonable any significant deviation from this as the default behavior of a broadband Internet service.

The IETF has not defined "Best Effort" in any Standards Track RFC; it is simply "understood" by the networking community[61]. This is unfortunate for the Commission's purposes; vague communal understanding is not a suitable basis for regulation. We tentatively offer the following definition, based on wording in RFC 2474:

> Best Effort Service: A network Service characterized by the following behaviors: When offered valid packets by user or edge providers, the network will deliver as many of these packets as possible and as soon as possible, subject to other resource policy constraints. When it must delay or drop packets during periods of congestion, the network will attempt to do so without unreasonable discrimination.

The precise definition of "fairness" is controversial in the networking community, and a number of different policies can be argued to be "fair". For example, the network might simply forward packets in strict FIFO order, and drop any that overflow the queue. Alternatively, it could service packet flows in a round-robin fashion, so as to prevent greedy flows from starving all others. Scheduling algorithms have unintended biases. Thus, instead of "fairness", we suggest an "unreasonable discrimination" test, which is more readily enforceable.

[60] RFC 7141 and RFC 2914 (together, BCP 41) describe best practices
[61] RFC 1812 describes forwarding steps associated with Best Effort service, but does not use that term.

We suggest that the Commission add a definition of "Best Effort Service" to the amended Rules at §8.11. We further suggest that the Rules explicitly state that special services might be offered in addition to the Best Effort default service.

A. Quantitative Performance Objectives

It would be a technical and procedural quagmire for the Commission to establish specific, quantitative edge-to-edge performance objectives for Best Effort service [62] in broadband networks. Any such objectives would have to be based on hard to justify assumptions, e.g., about a baseline work load and response time expectations. Lack of objective basis for these assumptions will lead to tussles among the experts. The assumptions will evolve over time, in unpredictable ways. For these reasons, we do not recommend this path.

A slightly different approach may be workable. The Rules could monitor and enforce minimum objectives for congestion of Best Effort traffic in network elements. For example, the objectives could specify maximum packet drop rate and maximum averaged queueing delay at each outgoing interface. Measurement could be at busy times of day, over medium term (e.g., 15 minute) intervals. This information exists in deployed network elements[63], and should be exposed. Chronic excessive dropping or delay for Best Effort traffic, while satisfying Specialized Services, would be presumptively unreasonable.

B. Reasonable Person Standard

A "reasonable person" standard for Best Effort service performance would seem to be an effective way of enforcing a No Blocking rule. Combined with enhanced disclosure rules as discussed above, it would give a basis for assessment of whether a broadband provider's offering is commercially reasonable. A network engineer, armed with the disclosures detailed above, could see whether a resource is congested for excessive periods, and whether that congestion is caused by excessive bandwidth reservation and/or an inadequate Best Effort pool. Such a condition would be presumptively unreasonable. Similarly, a classifier pattern match and discriminatory action against packets sent to or from an edge provider would be presumptively unreasonable; this would be subject to exceptions for network integrity, law enforcement, national security and explicit consumer request.

[62] e.g., for packet loss rate, delay variation, or throughput of TCP-compliant flows
[63] cite MIB

XII. BROADBAND PROVIDERS SHOULD BE PERMITTED TO OFFER SPECIALIZED SERVICES, SUBJECT TO COMPETITIVE AND CONSUMER SAFEGUARDS

We disagree vigorously with those commenters who would bar broadband providers from offering other services in addition to Best Effort[64]. At the same time, we recognize that without safeguards for consumers and edge providers, such services could be a breeding ground for the kinds of anti-competitive and monopoly rent seeking behaviors feared by the public interest parties.

To date, we have not seen Comments which provide strong technical justification for allowing specialized services. Most of the arguments we have seen are couched in terms of free market, anti-regulatory rhetoric[65]. We offer the following cases which might have technical and/or economic merit, yet could be claimed to violate strict Network Neutrality.

1. Best Effort service is dependent upon TCP slow start and congestion avoidance. This forms a closed-loop control system which attempts to regulate each flow to consume a "fair share" of the bandwidth of the bottleneck link. Some applications are more amenable to an open loop, rate based control system which allows transmission at or below a negotiated data rate without packet discard or excessive delay[66]. In particular, video and audio applications are better served by such services. "Over-the-top" video, voice and telepresence protocols and devices have evolved to "make do"[67]. However, in order to do so, they frequently suffer from significant buffering delays and degradations of video resolution and audio quality. Quality-of-Service guarantees would enhance user experience, potentially enough to want to pay for.
2. Broadband providers presently use specialized "managed" services to deliver their voice and video products. This validates their potential value to Over-the-Top competitors. It is ironic that in the interest of competition, public interest and other parties would have them deny these services to competitors.
3. Contrary to common perception, the Internet's congestion control system is not perfectly fair. This is an unintended consequence of the complex interaction of TCP's slow start and congestion avoidance mechanisms with packet discard behavior in congested routers. Most notably, the system is biased to favor established, high-rate flows of large packets, and starve new, lower rate

[64] cites
[65] e.g., "Micromanaging the Web Would Be a Macro Mistake", T. Hazlett and J. Wright, The Wall Street Journal, July 14, 2014. Also comments of American Enterprise Institute and AT&T.
[66] cite
[67] cite for VoIP, MPEG-DASH

flows and those with small packets. If applications that produce long-lived, high-rate flows are segregated within their own, separate control regime, performance of applications that produce short-lived and low-rate flows can improve.

4. Applications which emit solitary packets at infrequent intervals do not participate in the Internet's congestion control system. However, when packets must be dropped as a congestion signal, a router would drop these packets with the same probability as packets from long-lived flows. These sources do not create congestion, but are disciplined by the network as if they do. Many emerging "Internet of Things"/Machine-to-Machine (M2M) applications are in this category[68].
5. Best Effort service attempts to allocate congested resources in a nominally "fair" manner. Most resources are not congested most of the time. Some large file transfers, such as software updates, are not latency sensitive. Thus, they can be incented (e.g., by pricing) to defer to Best Effort traffic during periods of congestion. The networking research and engineering community has long discussed a "Bulk" service of this nature[69]. The IETF is encouraging experimentation along these lines[70].
6. The Internet's congestion control system depends on voluntary cooperation by user devices and edge equipment. Selfish behavior results in greater throughput for the offender, while reducing throughput for competing traffic. The Internet research community [71] has been exploring approaches to creating accountability for congestion. This could, for example, enable policing of misbehaving flows, or congestion pricing. It is expected to remove perverse disincentives to cooperative sharing of congested resources. The IETF is in early stages of standardizing such a scheme[72].

The research community finds the Internet's dynamics to be a fruitful field for their work. This is because of scaling due to growth, improved hardware technologies, changing application mix and lessons learned. We fear that if the Commission freezes the current Best Effort service model, many innovative research results will be precluded in the public Internet.

---

[68] RFC 5897

[69] RFC 6297 is a survey of research in this area.

[70] RFC 6817, RFC 3662

[71] e.g, Briscoe, Bob, et al. "Policing congestion response in an internetwork using re-feedback." ACM SIGCOMM Computer Communication Review 35.4 (2005): 277-288.

[72] Initially, RFC 6789. The CONEX Working Group has several drafts which are expected to become standards-track RFCs.

XIII. THE NPRM IS ON THE RIGHT TRACK IN CREATING AN ENFORCEABLE RULE

We do not presume to offer legal advice to the Commission. It appears that legal scholars disagree as to the authority granted under Sections 706 and 238 of the Act, even with the guidance of the *Verizon* court. We understand and concur with the Chairman's reasoning in trying to avoid reclassification of broadband providers under Title II of the Act, while maintaining a "nuclear option" to do so if absolutely necessary. We understand that Reclassification will be a heavy lift for the Commission, not only because of the inevitable backlash from broadband providers and the House of Representatives, but also because doing so will necessitate large numbers of forbearance proceedings. We would support doing so if policy goals we have recited are cannot be met under an enforceable Rule, or if a Rule grounded in the the Court's analysis cannot be sustained on appeal.

We do respond to some of the questions raised in the NPRM, from a technologist's perspective.

The Commission asks[73] for comments about a Rule "permit(ting) broadband providers to serve customers and carry traffic on an individually negotiated basis, "without having to hold themselves out to serve all comers indiscriminately on the same or standardized terms," so long as such conduct is commercially reasonable." Individually negotiated carriage, with secret pricing, terms and conditions, is extremely vulnerable to anti-competitive abuse. The standard of commercial reasonability must prohibit exclusionary contracts, self-dealing, predatory pricing, restraint of trade, and similar violations of anti-trust doctrine. The matter of pricing is of particular concern: preferential pricing of Specialized Services would greatly advantage one edge provider over another. Under the *Verizon* analysis, most forms of price regulation are not available to the Commission. We have no advice as to how a commercial reasonableness standard could prevent unfair pricing, but urge that it be formulated to do so.

The Commission asks for comments[74] as to the impact of individually negotiated services on edge providers who do not choose to negotiate. As we discussed above, the default service in the absence of negotiation is Best Effort. There is no impact, as long as the No Blocking Rule adequately protects the Best Effort service against resource starvation by individually negotiated services. The greater concern would be if the hypothetical entities – particularly the VoIP service provider - did wish to negotiate a better optimized service, but was rebuffed or offered disadvantageous pricing.

[73] NPRM at 116
[74] NPRM at 120

Also note that small political websites and narrowly focused social media would be unlikely to own their own infrastructure; doing so is economically impractical. Instead, they would deal with a hosting provider and/or CDN provider, who would in turn secure specialized services as needed. These entities would have greater negotiating power, and would not be subject to point-of-view discrimination.

The NPRM asks for comments on "concerns about the effect of pay-for-priority agreements on Internet openness"[75]. We believe that these concerns are adequately addressed by rules safeguarding the Best Effort service. Also note that simple "priority" is not at issue. We do not envision Edge Provider A paying to have their packets placed in private queues that are always serviced before queues of packets belonging to Edge Provider C. Nor would Edge Provider A's private queues be likely to be scheduled with strict priority over all best effort traffic. These behaviors would not be commercially reasonable. Edge Provider A might pay to have guaranteed timely delivery of packets conforming to a TSpec; in order to satisfy those guarantees, some of those packets might be scheduled ahead of some earlier Best Effort packets. Since individual packets do not get road rage when others 'cutting in front', it is sufficient that Best Effort traffic be assured sufficient resources under the No Blocking Rule.

The NPRM asks for comments concerning the application of a commercial reasonableness standard to vertical integrated broadband providers and affiliated entities[76]. In our view, this is the greatest competitive risk to the NPRM's framework. Scenarios such as a broadband provider exclusively providing special services to an affiliate, making exclusive contracts, and engaging in discriminatory pricing raise serious competitive concerns. The Rule must preclude these sorts of behaviors. While we cannot comment on its legal enforceability, "a rebuttable presumption that broadband provider conduct that forecloses rivals (of the provider or its affiliates) from the competing marketplace is commercially unreasonable" appears to delineate a minimum acceptable behavior. Pricing is a cause for special concern: how would the Commission or other affected parties detect pricing discrimination by a broadband provider in favor of an affiliated entity? The Rule may need to require disclosure of transfer pricing within a vertically integrated entity.

The NPRM asks for comments on a number of other points related to the commercial reasonableness standard. These points appear to be obvious, peripheral and/or redundant.

A. "Paid Prioritization" is not *per se* commercially unreasonable

[75] NPRM at 121, and footnote 250
[76] NPRM at 126-128

We strongly disagree with those who would claim that what they call "paid prioritization" is categorically commercially unreasonable. As we have discussed throughout these comments, specialized services, including those with usage-based or other charging models, can be permitted with competitive and consumer safeguards, such as the No Blocking Rule and the commercial unreasonableness standard. Such services, properly applied, can confer significant benefits to all parts of the Internet ecosystem.

That said, we would regard some forms of priority to be commercially unreasonable. These would include absolute priority implemented as to unacceptably degrade Best Effort service, and priority of one edge provider's specialized service over their rivals'. We will not comment on the legal issues concerning a list of practices that are *per se* unreasonable.

B. Safe Harbor for non-exclusive provision of special services to unaffiliated entities

AT&T suggests that non-exclusive agreements with unaffiliated edge providers be exempted from review[77]. We would prefer that specialized services be offered to all comers, at published terms, conditions and prices. This apparently would be dangerously close to common carriage, were the Rule to require it. However, there is no apparent prohibition against broadband providers doing so voluntarily. We agree with AT&T that there is no reason for the Commission and broadband providers to waste time and effort on duplicative case-by-case reviews of substantially identical agreements. However, we would limit the Safe Harbor to incremental agreements for special services that have been reviewed, and their terms, conditions, and prices disclosed to the public. Whether this would encroach on the *Verizon* guidelines is a question for legal experts.

XIV. THE NPRM APPEARS TO BE ON THE RIGHT TRACK IN ITS PROPOSED ENFORCEMENT AND DISPUTE RESOLUTION MECHANISMS

We concur with the Commission's approach to enforcement and dispute resolution[78]. Certainty, flexibility and accessibility are good goals. "Lightweight" procedures and structures like business letter reviews, non-binding staff opinions, enforcement advisories, an ombudsman, simplified procedural requirements, improved informal complaints processes, alternative dispute resolution, multi-stakeholder processes, and engagement of technical advisory groups all appear to be preferable to the Commission's existing "heavyweight" enforcement and dispute resolution processes.

[77] NPRM at 141
[78] NPRM at 161-176

In creating these mechanisms, and in organizing itself to execute them, the Commission must be cognizant of the fact that the matters to be adjudicated will be highly technical. This is esoteric technology, not widely understood except by experts. That this proceeding ran off the rails due to incorrect, imprecise and misleading framing of the technical issues should raise concerns. The Commission's Chief Technologist is a recognized and well respected expert in this technology. However, it appears – at least from the outside – that his depth of understanding does has not suffused the Commission and key staff. The Commission must develop a greater depth of technological understanding, through learning development, hiring and contractors.

XV. CONCLUSION

The Commission is generally on the right track. We have suggested a subtle reframing which we hope will alleviate some controversy, or at least ensure a defensible rationale.

At last count, almost 700,000 comments had been filed in this matter. We hope that these will stand out and be among those considered in developing the next Order.