

7-15-2014

# A “Most Favored Nation” Approach to an Open Internet

Jon M. Peha

*Carnegie Mellon University*, [peha@cmu.edu](mailto:peha@cmu.edu)

Follow this and additional works at: <http://repository.cmu.edu/epp>

 Part of the [Engineering Commons](#), and the [Public Policy Commons](#)

---

This Technical Report is brought to you for free and open access by the Carnegie Institute of Technology at Research Showcase @ CMU. It has been accepted for inclusion in Department of Engineering and Public Policy by an authorized administrator of Research Showcase @ CMU. For more information, please contact [research-showcase@andrew.cmu.edu](mailto:research-showcase@andrew.cmu.edu).

Before the  
FEDERAL COMMUNICATIONS COMMISSION  
Washington, DC

GN Docket No. 14-28

In the Matter of

Protecting and Promoting the Open Internet

A “Most Favored Nation” Approach to an Open Internet

Jon M. Peha

Professor, Carnegie Mellon University

Address: Carnegie Mellon University  
Department of EPP  
Pittsburgh, PA 15213-3890  
[peha@cmu.edu](mailto:peha@cmu.edu)

July 15, 2014

## Executive Summary

It is important for the Commission to establish new Open Internet rules. However, the proposed rules that set a minimum level of access and allow unlimited negotiation beyond that minimum will not protect the Open Internet. The NPRM's proposal to maintain the "specialized services" loophole is also a potential threat to an Open Internet. However, the NPRM is right not to prohibit all forms of priority, as priority can benefit Internet users. This paper proposes an alternative "most favored nation" approach, wherein arrangements between Internet access providers and content or application service providers must be public under transparency rules, and a content or application service provider is empowered under the "commercial reasonableness" standard to demand any arrangement that an Internet access provider has previously offered to another similarly situated entity. This policy would particularly benefit new entrepreneurial ventures. The Commission should simultaneously establish limits on the definition of "specialized services" to close loopholes in the 2010 rules that could allow the worst forms of "pay for priority," and the Commission should broaden its definition of "pay for priority." This most-favored-nation policy would protect the Open Internet far more effectively than the rules proposed in the NPRM, and the Commission has sufficient authority under either Section 706 or Title II. The Commission should also begin collecting more information on interconnection, without directly regulating the quality or price of interconnection.

## Author Qualifications and Disclaimer

Jon Peha is a Full Professor at Carnegie Mellon University, with experience in industry, government, and academia. In government, he served at the FCC as Chief Technologist, in the White House as Assistant Director of OSTP, in the House Energy & Commerce Committee, and at USAID for the Telecommunications Leadership Program. In industry, he has been Chief Technical Officer for three high-tech companies, and member of technical staff at SRI International, AT&T Bell Labs, and Microsoft. At Carnegie Mellon, he is a Professor in the Dept. of Electrical & Computer Engineering and the Dept. of Engineering & Public Policy, and former Associate Director of the Center for Wireless & Broadband Networking. Dr. Peha holds a PhD in electrical engineering from Stanford. He is an *IEEE Fellow* and an *AAAS Fellow*, and was selected by AAAS as one of 40 Featured Science and Technology Policy Fellows of the last 40 years ("40@40"). Dr. Peha has received the FCC's "Excellence in Engineering Award," the IEEE Communications Society TCCN Publication Award for career contributions, and the Brown Engineering Medal. He consults on a wide range of technical and policy issues related to information and communications technology.

In writing this comment, Dr. Peha represents no one but himself.

## Contents

1. [There is still need for Open Internet rules.](#)
2. [The FCC should focus more on how broadband providers can extract rents in upstream markets.](#)
3. [Some uses of priority benefit consumers and should not be prohibited.](#)
4. [A “minimum level of access” requirement will not protect the Open Internet.](#)
5. [If a minimum level of access is established, it should not be defined as proposed in the NPRM.](#)
6. [The FCC should close the “specialized service” loophole, which could allow the worst forms of pay for priority.](#)
7. [It is possible to prohibit the most harmful forms of discrimination by adding a “most favored nation” requirement to the “commercially reasonable” standard, and expanding the transparency requirement. The Commission has authority under either Section 706 or Title II.](#)
8. [The Commission should broaden its definition of “pay for priority.”](#)
9. [The FCC should collect more information about interconnection agreements and congestion.](#)

### 1. There is still need for Open Internet rules.

Network neutrality became an important Internet policy issue because of two trends: the emergence of traffic management and billing technologies that allowed Internet service providers to discriminate in sophisticated ways, and the emergence of a market structure that gave some Internet service providers incentive to use discrimination in a way that harms consumers and reduces social welfare.<sup>1</sup> The FCC seeks comment on whether the resulting need for Open Internet rules has changed since the rules were adopted.<sup>2</sup> As discussed below, each of these conditions remain, so the need for Open Internet rules remains as well.

The FCC and the D.C. Circuit court are both correct in their conclusion that “broadband providers have the technological ability to distinguish between and discriminate against certain types of Internet traffic.”<sup>3 4</sup> Technical methods that support this ability were in place before the Open Internet rules were adopted, and these same methods are still effective today. As described more fully elsewhere,<sup>5</sup> methods to distinguish between traffic types include inspection of network-layer and transport-layer *header information, deep packet inspection, and flow classification*. Methods to discriminate among traffic classes once traffic has been categorized include separation of traffic into *separate real or virtual channels*, and use of *traffic control algorithms* for functions such as packet scheduling, packet dropping, or routing that discriminate. All of these affect the quality of service that users observe. In addition,

---

<sup>1</sup> Section 1 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

<sup>2</sup> Paragraph 51 of Federal Communications Commission, Notice of Proposed Rulemaking in the matter of Protecting and Promoting the Open Internet, GN Docket No. 14-28, May 15, 2014. Henceforth Open Internet NPRM, May 2014.

<sup>3</sup> Paragraph 51 of Open Internet NPRM, May 2014.

<sup>4</sup> *Verizon v. FCC*, United States Court of Appeals for District of Columbia Circuit, No. 11-1355, Jan. 14, 2014.

<sup>5</sup> Section 2 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

Internet providers can use *content billing*, a technique by which users are charged based on the types of traffic they send and receive, independent of whether quality of service differs by traffic type. Similarly, they can charge users by the byte sent and/or received, but choose to count some but not all traffic in the byte count.

A noteworthy new development is the increasing number of direct peering agreements between Internet access providers and large content or application providers. For traffic associated with these agreements, an Internet access provider can often identify the application type (e.g. video streaming) or at least the specific provider of that content based on the point where traffic entered its network. Discrimination may occur at or near the Internet exchange point, or packets may be tagged when they enter to facilitate discrimination elsewhere, such as in the last-mile connection. As a result, some forms of discrimination are even easier today than they were a few years ago.

The other trend that made network neutrality important was the reduction of competition among Internet access providers when broadband replaced dial-up as the primary means of Internet access. In a highly competitive market, Internet access providers would be less likely to use discrimination in a way that end users dislike because some customers would switch to a competitor. However, if for example there is only one other Internet access provider in the region, and this competitor provides data rates that are vastly inferior, then an Internet access provider may not worry much about losing customers. We have not seen a substantial increase in the number of broadband Internet access providers available to end users since the Open Internet rules were first adopted.

## **2. The FCC should focus more on how broadband providers can extract rents in upstream markets.**

It is impossible to craft effective network neutrality rules without first understanding why profit-seeking broadband access providers might have incentive to adopt practices that harm consumers. For this reason, the NRPM's question on whether there are other economic theories that the Commission should consider<sup>6</sup> is especially important. As described in more detail elsewhere,<sup>7</sup> it is particularly important to consider rent extraction in upstream markets.

Competition among broadband access providers is limited, and limited competition can fuel rent extraction. Consider the case of a profit-seeking monopoly provider. In general, a monopoly maximizes profit by setting prices as high as possible without exceeding what consumers are willing to pay. If consumers are willing to pay more for product A than product B, then a monopoly will charge more for product A, even if costs are the same. Broadband Internet access providers sell the delivery of packets, and Internet users are willing to pay far more for the timely delivery of some packets than for others, depending on upstream market, where a market may be associated with a type of application, content

---

<sup>6</sup> Paragraph 50 of Open Internet NPRM, May 2014.

<sup>7</sup> Section 4 of J. M. Peha, "The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy," *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

or device attached at the customer premises. The profit-maximizing strategy is to charge differently for different upstream markets, such that prices reflect willingness to pay. For example, a broadband provider could increase profit by charging differently for the timely delivery of a video streaming packet than the timely delivery of a VOIP packet, and differently for the delivery of a 5 MB iTunes song than for the delivery of a 5 MB text file containing the latest FCC NPRM. One of the ways to achieve this is through “content billing.”<sup>8</sup> Note that in this way, broadband providers can extract rents even in highly competitive upstream markets. For example, even if there are hundreds of companies streaming music to consumers, a broadband access provider can add a charge for delivering this music that will drive total price towards what we expect in a music market with little or no competition. Because a policy that succeeds at maintaining competition in upstream markets can still fail at preventing oligopoly rent extraction in these markets, limiting the extraction of rents should be an explicit policy goal.

While content billing is not widely used in the United States today, there is a variant of this strategy that should fall within the definition of pay for priority that appears to be likely. For some applications, the value of delivering a given packet stream depends greatly on the quality of service, as reflected in end-to-end latency, packet loss rate, and related measures. For example, latency matters so much for a VOIP packet that delivering a VOIP packet too late is no better than blocking the packet entirely. For these applications, a monopoly broadband access provider can extract rents almost as effectively as described above by setting a price for good quality of service that depends on the upstream market, and implicitly on willingness to pay. As a result, we could see a different price for providing good quality of service for the streaming of the Super Bowl than for the streaming of an algebra tutorial for high school students, even though the broadband provider is providing the same service and incurring the same cost in both cases.<sup>9</sup> In a highly competitive market, rival broadband access providers might defeat this strategy by undercutting the more excessive prices, but this is less likely in markets with no more than two Internet access providers.

Thus, if a profit-seeking broadband access provider with sufficient market power can set different prices for each upstream market, this is likely to lead to oligopoly rents, higher prices for Internet users, and probably lower Internet utilization. Open Internet rules should address this scenario.

### **3. Some uses of priority benefit consumers and should not be prohibited.**

Although there are ways to employ priority in a way that extracts rents and harms Internet users, there are also ways to employ priority and other forms of discrimination that bring great value to Internet users.<sup>10</sup> The Commission should seek to preserve these valuable uses.

---

<sup>8</sup> Section 2 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

<sup>9</sup> The reader is free to judge which is the higher-value service.

<sup>10</sup> Section 3 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

One benefit of these mechanisms is to protect quality of service for those applications that need protection most. For example, videoconference applications would be far more valuable if it were possible to protect packets associated with these applications from long queueing delays even during periods of network congestion. If this led to a very slight increase in the time to download a large email attachment, many of us wouldn't mind or even notice. We may someday see life-or-death applications emerge that require better quality of service in the face of congestion. For example, researchers are advancing telemedicine applications for patients who connect from their homes to communicate with distant health care providers in real time. The FCC should not prohibit mechanisms that would make these applications practical.

Some have argued that Internet access providers should be allowed to prioritize some traffic, as long as the access providers do not charge for that service. However, differential pricing is intertwined with differential quality of service. Unless there is some form of higher price associated with superior service, we may see all users requesting superior service for all traffic, which would defeat the purpose.<sup>11</sup> Moreover, if there is a cost to provide these services, Internet service providers are less likely to provide the service if charges are prohibited.

Others argue that priority services should be allowed, and that Internet access providers should be able to charge for them, but Internet access providers should only be able to charge customers with direct connections rather than remote content or application service providers. This limitation has the benefit of making some forms of rent-seeking more difficult, but there are two important disadvantages. One concerns financial arrangements. If an Internet access provider is allowed to charge residential customers but not content providers, then it becomes more difficult to offer services in which residential Internet users are not expected to pay. For example, a university might wish to offer a massively open online course (MOOC) to high school students throughout the U.S. without requiring students to pay anything, including the cost of enhanced quality of service from an Internet access provider, but this would be prohibited. A for-profit video provider that gets its revenue from advertising would have a similar problem, as would its viewers who prefer not to pay for TV. The other disadvantage of this limitation is technical. The IP protocol makes it easy for Internet users to indicate in the packets they send that the sender wants these packets to experience superior quality of service and that the sender is willing to pay for this. However, there is at present no way for a user to indicate that there is a particular set of packets she wants to receive with superior quality of service, and that she commits to paying for this service. Although this problem is not insurmountable, it is not easy, and would require new technology and new business practices.

#### **4. A “minimum level of access” requirement will not protect the Open Internet.**

A “no blocking” rule and a transparency rule are important, but they are not enough. The recent NRPM proposes that the Commission impose a “minimum level of access” on broadband Internet access

---

<sup>11</sup> Section 3 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

providers, and then allow them to negotiate as they wish for service above this level.<sup>12</sup> As will be discussed below, this will not protect the Open Internet by preventing rent extraction, and it undermines the Commission's case to use Section 706 of the 1996 Telecom Act for legal authority.

If the Commission establishes a minimum level of access that is beyond what anyone would ever want, then no one would ever ask an Internet access provider for service exceeding the minimum. Thus, this is in effect equivalent to prohibiting such agreements, but with additional, unnecessary and possibly burdensome constraints on operations. If that is the Commission's intention, it should simply prohibit the arrangement outright. Thus, we assume here that there are some applications for which performance beyond the minimum is desired, no matter how that minimum is defined. Price and quality of service for these applications would be subject to open negotiation, and as some propose, unfettered negotiation as long as the deal is not entirely exclusive and not with an affiliate of the Internet provider.<sup>13</sup>

In this case, many forms of rent extraction discussed above are available to an Internet access provider with sufficient market power. In the example above, a provider might charge very differently for service that is better than the minimum when delivering the Super Bowl versus when delivering an algebra tutorial. More generally, an Internet access provider could set prices to extract oligopoly rents by exploiting the fact that user demand depends greatly on application and/or content. As discussed in Section 2, merely preventing an affiliate of the Internet access provider from monopolizing the upstream market does not prevent this form of rent extraction. An Internet access provider could also set prices in a way that protects legacy services from new over-the-top competition, perhaps by making over-the-top VOIP expensive. Equivalently, instead of setting different prices for the delivery of traffic at a given premium quality of service, a provider might hold price constant while offering different levels of quality of service to extract rents or to protect legacy services, perhaps by making sure that the quality of over-the-top VOIP is never as dependable as legacy telephone service.

The fact that this policy is relatively ineffective at protecting the open Internet also means it is less likely to withstand legal challenge if the Commission bases its authority on Section 706 rather than Title II of the Communications Act. The court was most persuaded by the argument that the FCC has the authority it claimed as long as Open Internet rules promote a virtuous cycle of innovation that leads to Internet expansion.<sup>14</sup> Consider an entrepreneur with a new application that would benefit from a quality of service beyond the minimum level. Not only must the entrepreneur pay for this premium service, but she may need to dispatch negotiators with extensive knowledge to every Internet access provider in the nation to get that service. Case-by-case negotiations can make transaction costs high, deterring this new entrant. From both a policy and a legal perspective, there are better approaches, as will be discussed in Section 7 of this document.

---

<sup>12</sup> Paragraph 95 of Open Internet NPRM, May 2014.

<sup>13</sup> Paragraph 141 of Open Internet NPRM, May 2014.

<sup>14</sup> *Verizon v. FCC*, United States Court of Appeals for District of Columbia Circuit, No. 11-1355, Jan. 14, 2014.



## **5. If a minimum level of access is established, it should not be defined as proposed in the NPRM.**

The NPRM proposes three methods of defining a minimum level of access. This section explains why all three are problematic, and presents an alternative.

One approach is to impose “a requirement that broadband providers apply no less than a ‘best effort’ standard to deliver traffic to end users.”<sup>15</sup> Best-effort delivery simply means that the network does not guarantee that packets will be delivered, so end systems that require reliable delivery may need to retransmit those packets that are lost. This proposal is unworkable for several reasons. First, it is meaningless to prohibit an Internet access provider from making its minimum-quality service worse than best effort because nothing is worse than best effort. What could be worse than unguaranteed? It is like prohibiting liquor stores from serving customers whose age is less than zero. Second, whether a service is best effort or guaranteed tells you little about what actually matters in this context - whether the quality of service provided is adequate for any given application. To someone watching a streaming video, quality depends on throughput. To someone on a VOIP call, quality depends on latency. There are best-effort services today that offer 100 Mb/s throughput, cross-country latency under 50 ms, and negligible loss rates. There are also best-effort services that provide 50 kb/s throughout, latencies of 250 ms, and high loss rate. We must look to something other than whether the service is best-effort to determine whether applications would be well supported. Third, the current Internet is based on IP technology, and IP is a best effort protocol. Typically, end devices use protocols such as TCP to make sure packets are delivered, but Internet service providers do not. Thus, until the dominant technology changes, it is likely that even premium services would be best effort.

Another proposed approach to “define a minimum level of access is through specific technical parameters, such as a minimum speed.”<sup>16</sup> As discussed above, application performance depends on parameters such as throughput, latency, and loss, so parameters such as these could form the basis of a more reasonable limit. However, there is no way to define parameters such as these in a way that applies to all technologies and situations in a useful and fair way. For example, consider one Internet access provider serving a remote village in rural Alaska where backhaul is scarce, and another serving Manhattan. Technical parameters that would seem reasonable for the Alaskan village would be so lax as to be irrelevant in Manhattan. Technical parameters that would seem reasonable in Manhattan would be so costly to meet in rural Alaska that they would in practice be unachievable. Thus, this approach may require defining different minimums for different systems, which would be difficult to do fairly.

The third proposed approach is to require Internet access providers to offer service at a “level that satisfies the reasonable expectations of a typical end user.”<sup>17</sup> The NPRM suggests that today a typical user could reasonably expect to be able to access streaming video, VOIP, and lawful content, but the

---

<sup>15</sup> Paragraph 102 of Open Internet NPRM, May 2014.

<sup>16</sup> Paragraph 103 of Open Internet NPRM, May 2014.

<sup>17</sup> Paragraph 104 of Open Internet NPRM, May 2014.

NPRM suggests nothing about how this was determined, or the quality of service currently deemed reasonable for these applications, or whether acceptable quality of service is the same in urban and rural areas, or how the reasonableness standard should be adjusted as technology and applications change. Without considerably more guidance, it is hard to see how the Commission would ever determine which set of user expectations are reasonable and which are not without constantly being challenged in court as acting in an arbitrary and capricious manner.

I will argue in Section 7 that there are better ways to protect an Open Internet than with a minimum level of access, but if a minimum level is needed, it would be better to base that minimum on how capacity is divided up. The advantage of this approach is that it can be applied both to a fiberoptic network capable of delivering 1 Gb/s service and a remote WISP that is constrained by its 10 Mb/s backhaul. For example, the Commission could say that no more than X% of Internet capacity in the more congestion-prone links of the network can be used for traffic associated with a premium service.

## **6. The FCC should close the “specialized service” loophole, which could allow the worst forms of pay for priority.**

While many believe that the 2010 Open Internet rules prohibited pay-for-priority arrangements, this may not be true. The 2010 Open Internet rules recognized that “specialized services” were not subject to all Open Internet rules, and the Commission declined to define specialized services. There are good reasons to label some services as specialized, and exempt them from some or all Open Internet rules. However, by not limiting use of the specialized label, the Commission gave Internet access providers the opportunity to adopt many practices that would otherwise be prohibited, including but not limited to pay-for-priority arrangements. The NPRM proposes to keep this potentially large loophole,<sup>18</sup> but this is the perfect time to limit what can be labeled as a specialized service.

Consider the case of an Internet access provider that also advertises a telephone service. The provider carves out a portion of IP capacity which is used exclusively for VOIP traffic associated with this telephone service. The service is regulated as a telephone service under Title II of the Communications Act. If this were an Internet service, it would violate Open Internet rules by blocking all applications other than VOIP. However, this service should clearly be treated as a telephone service that happens to be IP-based. This could reasonably be labeled as a specialized service.

Now the Internet access provider would like to add a specialized service for access to something that is a substitute for a service that is widely available on the Internet. For example, this specialized service is for auctions, and it looks a lot like eBay. (It may even be a repackaged eBay.) The IP-based traffic associated with this service shares capacity with other Internet traffic, but the auction traffic gets better quality of service than most Internet traffic. If no Open Internet rules apply, an arrangement like this could be used to give an affiliated auction site better quality of service than all of its competitors, or it could be used to make traffic from one upstream market more expensive than traffic from another

---

<sup>18</sup> Paragraph 60 of Open Internet NPRM, May 2014.

upstream market even if the cost of carrying the traffic is the same for both. In effect, this is a pay-for-priority arrangement with no limits on rent extraction. Thus, it should not be possible for Internet access providers to use the specialized services label in this case, or if they can, additional limits on specialized services are needed.

**7. It is possible to prohibit the most harmful forms of discrimination by adding a “most favored nation” requirement to the “commercially reasonable” standard, and expanding the transparency requirement. The Commission has authority under either Section 706 or Title II.**

The FCC could continue to allow pay for priority arrangements and still prohibit many forms of oligopoly rent-seeking by allowing content and application service providers to invoke “most favored nation” rights. This means that any content or application service provider can demand that an Internet access provider offer it the same deal that this Internet access provider has recently given some other similarly situated content or application service provider. This deal includes both the quality of service offered by the Internet provider, and the pricing. To grant this deal to one content or application service provider and then deny the deal to another similar situated content or application service provider would not be considered “commercially reasonable.” A “most favored nation” rule severely limits the extent to which an Internet access provider with market power can exploit differences in willingness to pay to extract rents. Such a network could at first serve the upstream markets with high demand at a high price, but once the network lowers price to serve other upstream markets with similar quality of service needs, the network would be prevented from raising prices again any time soon.

The Commission could adopt this approach whether it reclassifies Internet access as a Title II service or not. If the Commission chooses not to reclassify, the courts are more likely to find that Section 706 gives the FCC authority for the “most favored nation” approach than for the policy proposed in this NPRM. The DC Appeals Court has indicated that a policy allowing negotiation is distinct from common carrier regulation, and the “most favored nation” approach allows individual negotiation. Any time an Internet access provider and a content provider wish to adopt an agreement for paid priority that they deem to be of mutual benefit, the parties are entitled to do so, even if the terms of this agreement are unlike any that have come before. These arrangements are therefore made in a manner that is demonstrably different from what is allowed under common carrier regulation, as the Court requires. The “most favored nation” approach also facilitates the virtuous circle of innovation that the court found convincing. Consider that entrepreneur with a new idea for an application. The entrepreneur is free to negotiate directly with the Internet access provider to get better quality of service than is typically available, but this can be costly, and it isn’t necessary. Instead, the entrepreneur could examine the deals that have already been established by knowledgeable negotiators representing well-funded content and application service providers, and try to find the deal that will work best for this new application. This facilitates entry by small innovators.

The approach described above only works if content and application service providers know about the pay-for-priority arrangements of the past. Thus, the Commission should collect and publish information on those arrangements.<sup>19</sup>

## **8. The Commission should broaden its definition of “pay for priority.”**

The NPRM says that “pay-for-priority arrangements” are arrangements “in which the broadband provider would agree with a third party to directly or indirectly prioritize some traffic over other traffic to reach the provider’s subscribers.”<sup>20</sup> This could be interpreted as an inappropriately narrow definition. There are countless technical ways in which an Internet access provider can give some traffic classes better quality of service than others that might not meet this definition,<sup>21</sup> but could still be effective in extracting rents. For example, an Internet access provider might make sure that 10% of capacity on a bottleneck link is always available to the favored traffic class during periods when there are packets from that class queued, or the Internet access provider may give all classes of traffic the same quality of service but some classes count towards a customer’s monthly quota and others don’t. The Commission should make clear that these arrangements are included too. The real issue is discrimination, and priority as described in the NPRM only includes some of many ways of discriminating

## **9. The Commission should collect more information about interconnection agreements and congestion.**

Open Internet rules did not address interconnection by design. However, interconnection practices have been changing, and some have argued that Open Internet rules should be expanded to address interconnection.<sup>22</sup> Some point to recent peering disputes. The Commission should consider greater transparency in interconnection at this time, but not the regulation of quality or pricing. Thus, “no blocking” and “unreasonable discrimination” rules should not apply unless an agreement addresses the quality of service within the Internet access provider’s network, in addition to interconnection itself.

It is theoretically possible that an Internet provider could use interconnection agreements in a way that harms the Open Internet. For example, a monopoly Internet access provider could deliberately allow its connections to transit networks to become congested, and then demand extremely high prices for peering from content providers whose application is harmed by congestion, or even deny peering to content providers that compete with affiliates of the Internet access provider. If this occurred, it could undermine the virtuous cycle of innovation and expansion, because smaller providers with new applications cannot easily establish direct peering and are more dependent on transit. Thus, there is a strong argument that the FCC has authority under Section 706 to act if this is occurring, and there is

---

<sup>19</sup> Paragraph 88 of Open Internet NPRM, May 2014.

<sup>20</sup> Paragraph 36 of Open Internet NPRM, May 2014.

<sup>21</sup> Section 3 of J. M. Peha, “The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy,” *Telecom Policy Research Conference*, 2006. [www.ece.cmu.edu/~peha/balanced\\_net\\_neutrality\\_policy.pdf](http://www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf)

<sup>22</sup> Paragraph 59 of Open Internet NPRM, May 2014.

certainly a strong argument for authority under Title II of the Communications Act as well. However, that does not mean FCC intervention in interconnection agreements would be helpful. The mere existence of public disputes over peering is not evidence that anything like the scenario described above is occurring, or that it would be in any Internet access provider's interest to adopt such a strategy for rent extraction. After all, this strategy may increase an Internet access provider's revenues from peering, but it could decrease revenues from end users by making Internet access less valuable to them. The Commission needs more information than what is publicly available today, and from a wide range of providers.

Moreover, if more congestion information were publicly available, this alone may help prevent scenarios such as the one described above, in addition to helping Internet users choose the best service for their needs. End users who are paying for a high data-rate last-mile connection might be particularly unhappy to learn that the content they want most is unavailable at high speeds because of congestion at the interconnection points. This might lead these users to subscribe to a less expensive plan, or to switch providers if there is a viable competitor. Internet access providers have incentive to avoid both of these outcomes, so this could increase incentive to prevent congestion at Internet exchange points.