



Nick Feamster
Professor, Department of Computer Science
Deputy Director, Center for Information
Technology Policy

310 Sherrerd Hall
Center for Information Technology Policy
Princeton University
Princeton, New Jersey 08540-5233
+1 609 258 2203
feamster@cs.princeton.edu

March 12, 2018

To Whom It May Concern:

My name is Nick Feamster, and I am a Professor of Computer Science at Princeton University, and the Deputy Director of the Center for Information Technology Policy. My research involves the design, implementation, and deployment techniques for measuring the performance of Internet access. Since 2010, we have been designing techniques to measure the performance of broadband access networks. Many of the techniques that my research lab has designed have ultimately been incorporated in the Federal Communications Commission's "Measuring Broadband America" program. We also conducted the first independent analysis of the MBA program data in 2011, which was published in the premier computer networking conference.

I have read Assembly Bills No. 2131, 2132, and 2139 before the New Jersey Legislature, which seek to provide incentives to Internet service providers to implement several of the "bright line" rules from the FCC's Open Internet Order of 2015 and to be transparent about network management practices. Below, I offer technical commentary that I believe will be pertinent to today's discussion.

Let me start by saying that as a technologist, I favor simplicity of implementation. My research career has been devoted to developing technical solutions that are easy to implement and deploy. The bills before the assembly are sufficiently "light touch" that an Internet service provider without undue operational burden could implement them.

Other testimony that I have read for this hearing appears to discuss "heavy-handed" regulatory approaches taken at the federal level (such as Title II classification), as well as the authority of various federal agencies such as the Federal Trade Commission (FTC) to enforce the bright line rules that are enumerated in the Assembly Bills. I am not an expert in federal vs. state authority and cannot comment on the legal nuances.

However, from a technical perspective, my reading of the Assembly Bills is as follows:

- The bills encode requirements that prevent paid prioritization, guarantee access to all lawful Internet content, and prevent any practice that impairs or degrades access to lawful content ("no throttling" and "no blocking").
- At least one bill requires disclosure to customers of network management practices and performance. ("transparency")

I believe that the wording of these bills is sufficiently "light touch" from both a legal and operational perspective that it is not appropriate to cast them in the same vein as their federal counterparts. Let me start with the less controversial points concerning blocking, throttling, transparency, before turning to prioritization.

Access to Lawful Internet Content ("No Blocking", "No Throttling")

Consumers should enjoy unimpaired, un-degraded access to lawful Internet services. Attempts to access lawful services should not be blocked or throttled. This notion of "free and open access to the Internet" is what differentiates us from other authoritarian states, and these provisions should be completely uncontroversial.

From an operational standpoint, they also cost nothing to implement, since from an operational perspective, they amount to doing nothing. As a side note, the inability to throttle certain types of lawful traffic may of course impose additional traffic load on the network. There may be situations where “reasonable network management” involves prioritizing certain traffic over others, as I discuss below in the section on prioritization. Needless to say, the preferred solution should always be to add capacity, and prioritization for network management can and should be distinguished from paid prioritization.

Although it is not the topic of this hearing, it should be noted that, while we these bills concern the practices of Internet service providers, content and application providers—from commerce sites to search engines—are *regularly* blocking access to content and products, both as an anti-competitive practice or as a stipulation for being allowed to conduct business in an authoritarian country.

In short, “no blocking” and “no throttling” are extremely light touch—and ultimately we should carry these principles beyond just the Internet service providers, to application providers broadly defined.

Transparency

The FCC’s Restoring Internet Freedom Order (paragraphs 220 and following) also talk about the importance of transparency and disclosure, and that the order is in fact more prescriptive than the Assembly Bills. Therefore, the transparency requirements of the Assembly Bills should be relatively uncontroversial.

The main question, of course, is what should be disclosed to consumers. This question is challenging. It involves considering what would be both meaningful and informative to consumers, while also transparent about practices. Disclosing the parameters of a token bucket traffic shaper, for example, might be sufficiently transparent for some cases, but utterly confusing to consumers. On the other hand, a weather forecast icon (“sunny”, “partly cloudy with a chance of packet drops”) may be easy for a consumer to digest but so vague as to be meaningless. Striking the right balance is important.

We should also talk about transparency in a way that imposes a light touch, in an operational sense for ISPs. Below, I discuss various metrics that ISPs *already collect*, and could disclose to consumers without disclosing proprietary aspects of their network design or business models. In an ideal world, my “wish list” might be longer; in my suggestions below, I have tried to be cognizant of operational and business realities while still ensuring that the consumer can get meaningful information about the services they have purchased.

What should be disclosed to consumers? Bill No. 2139 in particular requires “public disclosure...regarding the network management practices and performance, and commercial terms of its Internet service”. Consumers need to understand the products and services that they are buying from their Internet service providers. The types of information about network performance that ISPs could reasonably provide consumers include:

- Upload and download throughput, as well as latency between customer and measurement server, as collected by the speedtest tools that they deploy (e.g., speedtest.xfinity.com), or that they contract to third parties to deploy (e.g., speedtest.net), across customers and regions.
- Periods of outage or service interruptions, affecting a significant number of customers (e.g., all customers associated with a cable modem head end) for an extended period of time.
- Latency statistics between vantage points in the ISP network and servers in popular application service providers (e.g., Netflix).

In addition to performance statistics that are useful for consumers, ISPs should also disclose reasonable network management practices. Note that reasonable network management practices **may include prioritizing certain types of traffic classes over others**—for example, all real-time interactive traffic

might be preferred over (say) “bulk” downloads or cloud backups. ISPs should disclose any such practices. Note that such prioritization is different than so-called “paid prioritization”, which I explain in more detail below.

Prioritization

Prioritization for network management vs. “paid prioritization”. All network traffic does not have the same requirements. Real-time or interactive applications such as video conferencing, the web, chat, or gaming have stricter latency requirements than a backup of a large file to a cloud service or the download of a large software update. Recognizing the differences in application requirements, it *can* make sense to prioritize certain types of application traffic over others, in a way that real-time application traffic can benefit from better performance.

In other words, slowing a software update by a small fraction so that users can enjoy the benefits of real-time applications can sometimes make sense. In this sense, it is important to draw distinctions between prioritization and *paid* prioritization.

Prioritization is a long-held network management practice that can achieve improved network performance for a class of network applications without unduly harming the performance of others. Prioritization might apply to *all* traffic of a certain type (e.g., all gaming traffic) as opposed to paid prioritization, which would give a higher priority to traffic only from one paying customer gaming site but not other competing sites.

In this light, one might view prioritization in the following light:

- When applied for the benefit of *all* traffic of a certain traffic class or application type—regardless of the customer application service provider—prioritization can be considered a reasonable network management practice.
- On the other hand, an ISP might apply prioritization to a set (or subset) of traffic from a particular customer as an anti-competitive practice, or in exchange for payment. For example, an ISP could (1) prioritize application traffic from one video service over another, if one service provider pays more money; (2) prioritize traffic for its own content over competing services.

Internet services may evolve—it would have been difficult to predict 15 years ago that the dominant traffic on the Internet would be video. Likewise, it is challenging to know for certain what the next 15 years will bring in terms of applications (and their requirements), and the laws must be flexible enough to account for this uncertainty. At the same time, the spirit of the prohibition on paid prioritization is to prevent anti-competitive practices, which is desirable.

The laws should allow for reasonable network management practices associated with prioritization, while prohibiting anti-competitive practices associated with paid prioritization.

Prioritization comes in many forms. In the vernacular, we often hear about an ISP “prioritizing” or “throttling” traffic from an application service provider. Yet, in the case of video traffic, the issue at hand is often a question of capacity provisioning—and who should pay for the provisioning of that capacity. Specifically, a large fraction of Internet video traffic is delivered via Content Delivery Networks (CDNs)—large networks of distributed servers that are often placed close to consumers to improve the delivery of application content such as video.

Traffic delivered from one application provider to the ISP’s customers may be delivered in a “best effort” (i.e., not prioritized) manner, but that application provider can still enjoy a distinct advantage over other application providers if it can gain access to (or deploy) a CDN that is close to users. When considering paid prioritization, it is important to also consider the indirect benefits and advantages that certain content providers may have over others, merely by virtue of the fact that they enjoy privileged access to servers

that are closer to consumers. In other words, an ISP may be **completely neutral** in how it treats the bits and flows from individual applications, but certain application providers may still enjoy advantages over others merely because of the access to physical infrastructure that they have.

Prioritization only matters if there is insufficient capacity. One of the first things that undergraduate computer science majors in a networking class learn is that common traffic schedulers give each application the rates that they ask for *as long as the capacity of the network exceeds the available demand*. As long as there is spare capacity, prioritization is a red herring because every application will achieve its desired sending rate. When capacity is scarce, however, prioritization becomes more important: a scheduling algorithm will determine how capacity should be allocated among the traffic flows, and some flows will not achieve the rates that they request.

Therefore, it is important to ensure that ISPs and application providers provision capacity in their networks accordingly, so that spare capacity exists. Spare capacity should exist both within a network, as well as on interconnects between ISPs and their neighboring networks (i.e., the application provider network or a transit network).

Summary

The Assembly Bills incentivize ISPs to eschew blocking, throttling, and paid prioritization, as well as to disclose various network management practices. The bills are written in such a way that they can be implemented without undue additional operational burden to ISPs. No blocking and no throttling requirements should be uncontroversial and are trivial to implement—this, as the ISPs themselves argue, is the status quo, in any case. In the case of the transparency requirements, the bills are even less constrained than the FCC's Restoring Internet Freedom Order.

Prioritization and paid prioritization are thornier, but there is a tractable approach here, too. When considering this topic, it is important to carve out reasonable network management practices from anti-competitive ones. In the presence of the transparency and disclosure rules that the Bills outline, it should be reasonable for ISPs to draw this distinction, and to disclose these any prioritization practices to consumers when it deems them necessary for operating the network.

Sincerely,

A handwritten signature in black ink, appearing to read "Nick Feamster". The signature is stylized with a large loop and a long horizontal stroke extending to the right.

Nick Feamster