A Free Press/New America Foundation Policy Brief:

The Hidden Harms of Application Bias

M. Chris Riley[1]
Free Press

Robb Topolski[2]
New America Foundation, Open Technology Initiative

November, 2009

Abstract

Application bias, the practice by network operators of placing applications into tiers of low and high priority and enforcing this prioritization through mechanisms in the middle of the network, poses hidden harms for the Internet that substantially outweigh its uncertain benefits. Application bias degrades low priority applications, decreases overall network performance, and locks the Internet into typical usage patterns of 2009, frustrating both consumer choice and Internet innovation. At the same time, the biggest hurdle to offering more powerful services on the Internet is not congestion, but rather delayed deployment of truly high-speed Internet access services. Application bias cannot make a small pipe into a larger one—prioritization cannot turn a rural DSL line into a fiberoptic cable, though it can perversely create disincentives from doing just that. With application bias, the benefits remain uncertain, but the risks to a dynamic and diverse Internet future are crystal clear.

[1] Chris Riley is Policy Counsel at Free Press. Dr. Riley holds a Ph.D. in Computer Science from Johns Hopkins University and a J.D. from Yale Law School.

[2] Robb Topolski is Chief Technologist at the Open Technology Initiative at New America Foundation. Mr. Topolski has over 25 years of experience with digital communications in wired and wireless networks.

INTRODUCTION

Central to net neutrality is nondiscrimination—the principle that content and applications on the Internet should not be treated differently by the network operator. As the FCC receives comment on its notice of proposed rulemaking to protect the open Internet,[3] the question of nondiscrimination will undoubtedly take center stage. One debated aspect of this question concerns whether to permit *application bias*—network management that places applications into tiers, giving the higher tier explicit priority over the lower. Application bias occurs when network operators choose which applications or classes of applications are prioritized over others.[4]

Prioritization operates at the level of the network router, generally by delaying or dropping lower-priority packets in times of congestion.[5] Without congestion, all packets are forwarded as they arrive, without delays or drops. But with congestion, prioritization forwards higher priority packets ahead of other traffic, and lower priority packets are negatively affected until there are no higher priority packets to send. Prioritization operates by degrading and harming lower priority traffic, because (by definition) more low priority packets are delayed or dropped. This harm will be further exacerbated by attempts to circumvent prioritization, software workarounds that masquerade low priority traffic as high priority, triggering an inefficient and expensive arms race between developers and ISPs seeking to deploy ever more intrusive and computationally intensive inspection methods.[6]

As a result, the question is not *whether* prioritization may cause harm, because it will—the question is whether that tradeoff is worthwhile. Supporters of application bias contend that the benefit of the prioritized uses outweighs the detriment to others, although no detailed study of the impact of application bias in a variety of network environments has yet been conducted.

---

[3] *Preserving the Open Internet; Broadband Industry Practices*, GN Docket No. 09-191, WC Docket No. 07-52, Notice of Proposed Rulemaking, FCC 09-93 (rel. Oct. 22, 2009).

[4] When users have complete control over their own application priority, and when their choice does not impact the network connection of others (for example, prioritization applied solely at the end nodes, rather than in the middle of the network), prioritization can be beneficial without raising substantial harms.

[5] Early specifications of Internet Protocol included an "IP Precedence" field, which allowed packets to be labeled at varying levels of priority. *See, e.g.*, J. Postel (ed.), Information Sciences Institute, "Internet Protocol DARPA Internet Program Protocol Specification," RFC 791 at 3.1, IETF (Sept. 1981). Although the specification does not identify a unique method of enforcing the specified preference, it states that this is done "generally by accepting only traffic above a certain precedence at a time of high load." *Id.* Although this field is now rarely used according to its initial specification, other uses have developed that implement the same concept. *See, e.g.*, K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the differentiated services field (DS field) in the IPv4 and IPv6 headers," RFC 2474 at 4.2.1, IETF (Dec. 1998) (discussing the IP Precedence concept and history).

[6] The ultimate result of such an arms race is likely to be near-universal encryption, rendering DPI and many expensive router deployments (whether used for traffic shaping or for law enforcement or defense purposes) nearly useless.

At first blush, the tradeoff may sound reasonable. If one application or class of applications can tolerate additional delay better than another, an application bias—particularly one without any obvious anti-competitive motivations—seems beneficial.[7] A network management practice that could differentiate among applications and provide an improved experience for more latency-sensitive applications without significant negative consequences would be welcome; unfortunately, no such magic bullet exists. Application bias has at least three substantial negative consequences: It restricts consumer choice by effectively degrading the performance of some high-priority applications; it impedes innovation at the edges of the network; and it degrades overall network performance. At the same time, prioritization has not yet been demonstrated to be necessary, or substantially helpful, for dealing with problems caused by congestion.[8] The Internet developed and flourished without network operator prioritization.[9] Given the potential for harm, a clearer need should be demonstrated before changing the Internet's status quo.

Further, even where a need for congestion management is demonstrated, user-based congestion management—allowing users to specify the traffic they want prioritized, and throttling heavy users where necessary—offers the same suggested benefits as application bias, but with fewer harms. The concept of "reasonable network management" should therefore only consider application bias to be "reasonable" if the network operator can demonstrate the existence of a problem requiring management, and can show that less harmful remedies such as additional investment or user-based congestion management are not feasible or appropriate.

[7] The "without any obvious anti-competitive motivations" concept is charitably assumed in this paper, but in practice, these motivations are strong and should be taken into consideration when evaluating the merits and dangers of application bias. Application bias creates a substantial potential for anti-competitive behavior by creating artificial distinctions between competing applications—such as using one priority tier for some video, like RealPlayer, but a different tier for another, like YouTube. Distinguishing competing applications and granting a priority to one, especially if that priority can be bought or if it can be offered for free to a partner or an offering of the ISP, creates substantial distortions in the market for applications, introducing a large new class of harms above and beyond those emphasized in this paper.

[8] The potential benefits of deprioritization are clearer, particularly when performed by end-nodes by scaling back transmission speed at the first signs of congestion; these acts can reduce the total amount of data hitting the router, thus reducing congestion and improving performance for other traffic. *See, e.g.*, T. Tsugawa, G. Hasegawa, and M. Murata, "Background TCP data transfer with inline network measurement," in *Proceedings of APCC 2005* (Oct. 2005). According to some reports, the newest version of the BitTorrent protocol, uTorrent 2.0, will include this mechanism. Ernesto, "uTorrent 2.0 To Eliminate The Need for ISP Throttling," TorrentFreak, *at* http://torrentfreak.com/utorrent-2-0-to-elimininate-the-need-for-isp-throttling-091031/ (Oct. 31, 2009). On the other hand, since prioritizing traffic in the middle of the network does not slow the flow of traffic hitting the router (as flow volume is determined by end nodes), it does not lessen total congestion.

[9] *See, e.g.*, Jon M. Peha, "The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy," 34th Research Conference on Communication, Information, and Internet Policy, George Mason University, at 3 (2006) ("In a packet-switched network such as the Internet, information is sent through the network one packet at a time.... Traditionally, Internet packets were sent with equal priority and "best effort," i.e. with no guarantee of delivery.").

APPLICATION BIAS CAN ACTUALLY DECREASE NETWORK PERFORMANCE

Application bias is sometimes portrayed as efficient, which is misleading at best if not flatly incorrect, as prioritization decreases overall network performance. Multiple distinct elements of application bias diminish performance. First, deep packet inspection (DPI) would typically perform the work of application identification needed to enable application bias.[10] DPI may add significant delay to the total routing time for a packet,[11] and the longer a packet sits in the network, the more it contributes to congestion and end-to-end delay.[12] Although DPI tools are growing faster and more powerful, minimizing these delays, it will always be more time consuming to inspect a packet's body than to read and route based solely on the packet headers, because the original design of the IP packet set specific formats and values for IP headers to ease analysis.[13]

Second, prioritization produces more dropped and delayed packets, causing more packet retransmission, and thus requiring greater traffic to perform the same communication. The Internet generally operates along a First-In, First-Out (FIFO) model, in which packets are forwarded in the order in which they are received.[14] Prioritization changes this design by increasing the delay of some packets while reducing the delay of others. Therefore, even if prioritization keeps the average packet latency the same (i.e., if DPI does not increase average latency), it will still increase the standard deviation by slowing low priority packets and speeding up high priority. When packets are delayed too long or when they are dropped, the packet sender will often resend the packet.[15]

---

[10] Deep packet inspection or some other computationally intensive method must be used, because IP packet headers contain little useful information for traffic discrimination. *See* Peha, *supra* note 9, at 3-5 (identifying flow classification and deep packet inspection as the two primary methods of learning more than IP header information about network traffic, and characterizing both as "requir[ing] a great deal of processing"). The technology companies that manufacture and sell DPI and related equipment go so far as to advertise its capability for exactly this purpose. *See* M. Chris Riley and Ben Scott, "Deep Packet Inspection: The End of the Internet as We Know It?," Free Press (Mar. 2009), *at* http://www.freepress.net/files/Deep_Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf.

[11] *E.g.* Tom Nolle, "Deep packet inspection: Controversial but valuable traffic management tool," Search Telecom (July 27, 2009), *at* http://searchtelecom.techtarget.com/tip/0,289483,sid103_gci1362852,00.html ("The technical issues of deep packet inspection are as complex as the policy issues. The problem can be stated in one word: latency. Packet analysis takes time, and that increases handling delay at the point where the analysis takes place. The more points requiring traffic analysis, the more the delay. And the more sophisticated the analysis, the more the delay.").

[12] Congestion occurs when packets are received by a router faster than it can forward them; increased time to process packets will reduce the rate at which packets that can be forwarded by the router, thus causing even more packets to build up, and further increasing latency.

[13] *See, e.g.*, J. Postel (ed.), Information Sciences Institute, "Internet Protocol DARPA Internet Program Protocol Specification," RFC 791 at 3.1, IETF (Sept. 1981).

[14] *See, e.g.*, Fei Li, "Competitive Scheduling of Packets with Hard Deadlines in a Finite Capacity Queue," Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM), p. 1062–1070 (2009) ("Currently, most [Internet] routers forward packets in a First-In First-Out (FIFO) manner and treat all packets equally.").

[15] The specification for all TCP packets establishes this as a fundamental component of ensuring reliable transfer over an unreliable network. J. Postel (ed.), Information Sciences Institute, "Transmission Control Protocol DARPA Internet Program Protocol Specification," RFC 793 at 2.6 ("When the TCP transmits a segment containing data, it puts a copy on a retransmission queue and starts a timer; when the acknowledgment for that data is received, the segment is deleted from the queue. If the acknowledgment is not received before the timer runs out, the segment is retransmitted.").

Thus, increasing standard deviation for the same average latency will result in more retransmissions, causing more packets to traverse network routes multiple times, creating additional packet load and therefore additional congestion in the network.[16] Widely fluctuating latencies also violate a programming assumption of reasonably consistent latency and jitter—with high but consistent latency, most applications such as video streaming can function quite well by varying quality and buffer size, but inconsistent latency can throw off these techniques and produce a worse user experience than a slower but consistent connection would have otherwise supported.

Imposing application bias against the choice of users will trigger an arms race with users and developers whose application performance is negatively impacted by the technique. Technological tools allow Internet users and application developers to encrypt and obfuscate their traffic, allowing them to avoid detection by DPI tools or even to generate false positives as video or some other prioritized traffic. Internet operators would need to engage in more intensive flow classification methods to circumvent this, further increasing processing delays. These ever more powerful electronics would also come at substantial recurring cost—a series of upgrades expending resources that could otherwise have been spent increasing capacity. Increased use of encryption would also hinder the work of law enforcement and security agencies.[17]

Although prioritization may enable some packets to reach their destinations faster than they otherwise would, it does so at the expense of the non-prioritized traffic.[18] Since prioritization is only useful in congested networks, the non-prioritized traffic—already receiving poor performance because of congestion—may face extraordinarily poor performance. All told, prioritization in the network imposes substantial burdens on non-prioritized traffic; may improve some traffic somewhat during times of notable congestion; and reduces performance of the overall network. Increasing total network capacity, on the other hand, improves congestion, improves overall network performance even in times without congestion, and improves performance for all traffic.[19]

---

[16] If a TCP packet is retransmitted from its origin, it will travel again over any network edges the original packet traveled.

[17] *See, e.g.*, Patrick Foster, "MI5 comes out against cutting off Internet pirates," *Times Online* (Oct. 23, 2009), *at* http://www.timesonline.co.uk/tol/news/uk/crime/article6885923.ece (discussing statements by British police and intelligence services to oppose three-strikes laws, out of concern that the resulting incentives to encrypt traffic would render prosecution more difficult).

[18] Because prioritization causes some packets to be passed through congestion points where others would have been, those other packets are left longer in the queue, and are either delayed or dropped depending on the duration of congestion and the number of higher priority packets that need to be passed through.

[19] In fact, as net neutrality proponent Google has noted, "QoS quickly can become an unspoken rationale to maintain artificial broadband scarcity." Comments of Google, Inc, In the Matter of *Broadband Industry Practices*, WC Docket No. 07-52, p. 27-28 (2007).

APPLICATION BIAS LOCKS THE INTERNET INTO TYPICAL USAGE PATTERNS OF 2009, FRUSTRATING BOTH CONSUMERS AND INTERNET INNOVATORS

Application bias restricts consumer choice by limiting the inherent freedom of Internet users to choose different applications, and to choose to use them in different ways. For example, many peer-to-peer (P2P) file transfers may indeed be low priority and tolerant of significant delay—but P2P applications are also used for high priority tasks, including VoIP and video streaming applications,[20] along with esoteric but important uses such as transmitting satellite imagery from NASA to scientific researchers.[21] P2P applications in particular have a wide array of valuable but little known (or not yet even imagined) uses, all of which will be harmed if broad classifications of "low priority" are imposed. Users who choose to use P2P for these tasks may face substantial performance problems because of the inability of network operators to effectively determine when P2P is used for a high priority task.[22] A similar problem will likely impact users of virtual private networks (VPNs), as communications through a VPN may be either "high" or "low" priority, and a network operator is poorly positioned to tell the difference; as a result, VPN traffic will likely be classified as entirely high or low priority, regardless of its actual use, based on whether it is more typically high or low priority. Hard-coding network management based on "typical" usage thus subjects all Internet users to the tyranny of the majority, and introduces substantial errors into the correct classifications of priority. Users who "think different," using applications for new or existing but uncommon purposes, will be left by the wayside, fundamentally weakening the Internet's generativity.[23]

Application bias hurts innovation as well as user choice. Without application bias, the performance of an application or service online as compared to its competitors is determined by the design and engineering of the application or service. As a result, a new application with a better design or better engineering can enter a market and succeed against well-established incumbents.[24] Application bias uproots this dynamic, creating a

---

[20] Notably, CNN used a peer-to-peer plug-in to handle web streaming of the inauguration of President Obama. Janko Roettgers, "CNN: Inauguration P2P Stream a Success, Despite Backlash," *NewTeeVee*, *at* http://newteevee.com/2009/02/07/cnn-inauguration-p2p-stream-a-success-despite-backlash/ (Feb. 7, 2009). CNN parent company Turner Broadcast Systems chose to use P2P in the belief that P2P was the only technology capable of handling the expected massive loads. *Id.* Skype uses P2P as a fundamental component of its popular VoIP software. "P2P Telephony Explained – For Geeks Only," Skype, *at* http://www.skype.com/help/guides/p2pexplained/.

[21] Recognizing the technique's superiority for transferring large files, NASA has used BitTorrent to distribute satellite imagery from its "Visible Earth" program since early 2006. *See, e.g.*, "NASA Propels Downloads With BitTorrent," Wired Blogs (Jan. 27, 2006), *available at* http://www.webmonkey.com/blog/NASA_Propels_Downloads_with_BitTorrent.

[22] ISPs are not well positioned to identify which uses of P2P are for VoIP or streaming video or other high priority communications; as a result, all or most P2P uses would likely be classified as "low priority." Furthermore, any activity by ISPs to impose prioritization will likely drive users to encryption and obfuscation techniques, rendering accurate identification even more difficult.

[23] "Generativity," according to Professor Jonathan Zittrain, is "a system's capacity to produce unanticipated change through unfiltered contributions from broad and varied audiences." *See, e.g.*, Nate Anderson, "Book Review: Jonathan Zittrain's 'The Future of the Internet—And How to Stop It'," *Ars Technica*, *at* http://arstechnica.com/old/content/2008/06/book-review-2008-06-2-admin.ars (June 18, 2008) (citing JONATHAN ZITTRAIN, THE FUTURE OF THE INTERNET—AND HOW TO STOP IT (2008)).

[24] The "garage inventor" motif characterizing the Internet is well established. *See, e.g.*, Chairman Julius Genachowski, "The Open Internet: Preserving the Freedom to Innovate," The White House Blog, *at* http://www.whitehouse.gov/blog/The-Open-Internet-Preserving-the-Freedom-to-Innovate/ (Sept. 21, 2009)

world where incumbents' applications or classes of applications receive priority over new entrants, if those new entrants fit outside the defined incumbent classes of priority while still competing with the incumbent in end-user functionality.

Ultimately, any time application bias is permitted, the determinations of "high" and "low" priority will reflect typical Internet usage as of 2009. No engineer or policy maker can predict the future of innovation online, or even typical patterns of Internet usage in 2014.[25] And yet, application bias locks Internet users into 2009 usage patterns by creating strong performance incentives to use the Internet on terms pre-defined by the network operator. Creativity, diversity, and innovation suffer as a result.

A hypothetical may help illustrate these problems: Imagine if RealVideo, the video format used in RealPlayer, was classified as a priority application upon its original release in 1997. Upon its introduction in 2005, YouTube might not have received the same level of priority, because it uses a fundamentally different protocol and business model—YouTube hosts video itself, whereas RealVideo is hosted on individual websites. After eight years of prioritized use, the video quality of RealVideo would have held a substantial advantage over the new entrant, YouTube, which will be effectively degraded by the imposition of priority for RealPlayer. YouTube would have faced an uphill battle to adoption, being required to compete as a video service without priority; it might well have failed, while on a level playing field, it flourished.

APPLICATION BIAS HAS NOT BEEN SHOWN TO BE NECESSARY OR EVEN SUBSTANTIALLY BENEFICIAL

Many proponents of application bias assert that the technique is necessary for latency-sensitive applications like VoIP or video streaming. But even without prioritization, services like Skype and YouTube function rather well on the open Internet—in fact, Skype and Google (the owner of YouTube) are both strong net neutrality supporters. Separately, some parties say that prioritization will enable new services such as telemedicine. Although very high capacity services that require substantial upload and download speeds may not be feasible over the average residential Internet access service in the United States, the supposed ability of prioritization (and not any other form of congestion management) to correct these limitations has never been demonstrated.

Limitations on the functionality of latency-sensitive or high-bandwidth applications can come from two separate sources: limited maximum capacity (particularly in residential networks and in the upload direction), or congestion. Application bias (and, in fact, all other forms of prioritization) cannot provide *any benefit* for the former; the only remedy is to build a bigger pipe.[26] Application bias can help one application in dealing with the

("With the help of all stakeholders, the FCC can help secure a bright future for the Internet, and make sure that the garage, the basement, and the dorm room remain places where inventors can not only dream, but bring their ideas to life.").

[25] The past five years have seen tremendous emergence of new, popular uses of the Internet. To take three major examples: Facebook was created in 2004; YouTube in 2005; and Twitter in 2006. Social networks, Internet video, and microblogging are very popular now, but the primary portals for these were nonexistent or in their infancy five years ago.

[26] Prioritization generally operates by forwarding only higher priority traffic at times of congestion, *see* supra note 5. Even if packets are reordered or forwarded through a round-robin scheme or some other mechanism,

problems of congestion, by pushing other users and other applications to the side; however, there are nondiscriminatory methods of dealing with congestion that can confer the same benefits without the same harms. No study or analysis has yet demonstrated a *single application* which can work because of an application bias, but *cannot* work on the same Internet connection with other forms of congestion management.

USER-BASED CONGESTION MANAGEMENT IS EMINENTLY MORE REASONABLE THAN APPLICATION BIAS

Congestion on the Internet is not a new phenomenon; network operators have dealt with the Internet's growth spurts over history. Historically, congestion was managed through moderate subscriber speeds and limits on the number of subscribers sharing a single access connection. Subscribers do not all use their full Internet access speed at the same time, so service providers physically oversubscribe the network—an ISP can (and does) sell more shares of bandwidth than it has physical capacity to support. As the popularity of the Internet has increased and the demands for faster speeds and greater consumption by each user have increased, building ahead of growing demands in order to avoid congestion has become a normal part of operating an Internet business.

But now, in the United States, oversubscribed and underbuilt networks are the norm, in which far more capacity is promised than can be effectively supported. As a result, some peak-time congestion will occur. And, in shared last-mile networks such as cable and wireless, a few heavy users can contribute disproportionately to network congestion, leaving fewer resources for other users (although any attribution of blame to these few users must recognize that they are merely using the capacity sold to them by their service provider, and the constraints and congestion come from an underbuilt network with overpromised capacity, not from any illicit user behavior). Although the best way to improve congestion is to invest and increase capacity,[27] the most reasonable short-term way is to maintain users' control over their own traffic, and throttle specific heavy users (when absolutely necessary), not specific uses. User-based congestion management is superior to use-based congestion management in several ways. Use-based management harms competition and innovation by placing a thumb on the scale of choice of application, often in favor of existing applications and against new and emerging applications; in contrast, user-based management harms neither. Use-based management encourages the creation of scarcity, particularly if application providers can pay for access to the higher tier; user-based management, on the other hand, encourages investment in the network, to provide more capacity to all users. Use-based management encourages users to use other applications but does not encourage them to reduce their overall usage; user-based management encourages users to monitor and manage their bandwidth usage to maximize network efficiency for their needs, and thus has a greater positive impact on overall network congestion.

---

prioritization cannot forward packets faster than a best-efforts Internet router can forward them (because it only serves to select which packets to forward), and thus cannot increase maximum capacity.

[27] The Canadian government agency with jurisdiction over Internet Service Providers has recently established this as a core principle of regulatory policy, recognizing that investment should be the "primary solution" for congestion management, with some traffic management practices as a secondary option. *See* Canadian Radio-television and Telecommunications Commission, Telecom Regulatory Policy CRTC 2009-657: Review of the Internet traffic management practices of Internet service providers at 2 (Oct. 21, 2009), *at* http://www.crtc.gc.ca/eng/archive/2009/2009-657.pdf ("Network investment is a fundamental tool for dealing with network congestion and should continue to be the primary solution that ISPs use; however, investment alone does not obviate the need for certain [Internet traffic management practices].").

Ultimately, imposing an application bias pressures Internet users toward some uses of the Internet and away from others, restricting innovation and user choice and further locking in the Internet of 2009 and delaying the evolution and emergence of the Internet of 2019. There are other ways to deal with congestion, better ways that pose fewer dangers for competition, choice, and innovation.

NONDISCRIMINATION POLICY SHOULD PROHIBIT APPLICATION BIAS, SUBJECT ONLY TO THE NARROWEST CLASS OF REASONABLE NETWORK MANAGEMENT

The right solution is a strict rule against discrimination for applications or classes of applications, whether that discrimination takes the form of degradation or priority. Any other policy limits consumer choice and generates anti-competition and anti-innovation incentives, without conferring any clear and demonstrable benefit. If a network management practice violates this policy, the network operator should face a heavy burden to demonstrate that the practice is necessary in the circumstances of its use. And no form of application or service provider payment for priority should ever be considered reasonable, as this would eliminate user choice, strangle innovation, and create strong incentives to maintain scarcity and delay investment, without any technical justification.

In fact, a similar policy framework was recently established in law by the Canadian Radio-television and Telecommunications Commission (CRTC).[28] Under Canadian law, whenever a complaint is filed against a traffic management practice that "results in any degree of discrimination or preference," the Internet Service Provider must show that the practice is designed solely to meet a demonstrable need and nothing else, that the practice "results in discrimination or preference as little as reasonably possible," and that "network investment or other economic approaches alone would not reasonably address the need and effectively achieve the same purpose."[29]

A framework such as this can permit either prioritization or degradation, but only where it is truly needed. For example, the service provider would need to demonstrate that an application bias only occurs in rare circumstances of heavy congestion (the only setting in which application bias can confer any benefits), and only where the congestion is sufficient to create enough latency or jitter to negatively impact the quality of the application receiving prioritization.[30] The service provider would also need to show that nondiscriminatory network management could not confer the same benefits. Finally, the service provider would need to publicly disclose its methods, and subject them to peer review and approval by an independent panel of experts. Given its potential for harm compared to the success of the Internet thus far, a service provider wishing to implement an application bias should face a high burden of proof to show it is needed.

---

[28] *Id.*

[29] *Id.* at para. 43.

[30] Standard thresholds to ensure high quality voice and video are a latency of at most 150ms, and jitter of 30ms. *See, e.g.*, Chris Lewis & Steve Pickavance, "Implementing Quality of Service Over Cisco MPLS VPNs," *Cisco Press*, *available at* http://www.ciscopress.com/articles/article.asp?p=471096&seqNum=6.

CONCLUSION

Giving a priority to a few popular uses of the Internet through application bias imposes top-down control over consumer choice; even if it were a beneficial dictatorship, in which the ISP has no anti-competitive motive but merely seeks to improve the quality of the most popular and typically high priority uses, application bias is at best tyranny by the majority. The most common applications may work a little better, but uncommon or emerging uses of the Internet would be degraded, potentially to the point of being inoperative. This would undermine the Internet's core value as a free and open market for all voices and all competitors, as well as its tremendous innovative and generative potential. The Internet would no longer be a general purpose technology, but would instead become a specialized media delivery mechanism, optimized for the Internet of 2009. But the Internet and its most common uses in 2014 may be very different—assuming network operators leave it free and open, and ready for innovative and generative future use.

Application bias is not necessary, and does not create the right incentives for network operators, application developers, or users. Enforcing discriminatory treatment in the middle of the network—looking into the application layer to impose an artificial application bias onto packet routing—has significant negative consequences, ranging from network performance issues to more subtle harms to innovation and competition online. Application bias would effectively create two separate Internets, the Fast Internet for incumbents and the Slow Internet for smaller (and inherently less common) individual content producers, small businesses, and new entrants, undermining competition and innovation. Slow Internet applications and content would be unable to compete with those on the Fast Internet—separate, but definitely not equal.