INITIAL COMMENTS OF THE

OPEN INTERNET COALITION

TELECOM PUBLIC NOTICE CRTC 2008-19:

REVIEW OF THE INTERNET TRAFFIC MANAGEMENT PRACTICES OF INTERNET SERVICE PROVIDERS

FEBRUARY 23, 2009

EXECUTIVE SUMMARY

1. The Internet supercharges and combines the power of free speech and free markets, producing the greatest engine of democratic deliberation and open commerce since the introduction of the moveable-type press. Yet, the continued success of the Interent is not inevitable and Canada's knowledge economy depends on the CRTC protecting the open Internet and facilitiating its continued growth.

2. Certain traffic management practices by Canadian carrier Internet service providers ("ISPs") threaten the open and neutral design of the Internet. The Commission should prohibit such practices as contrary to sections 27(2) and 36 of the *Telecommunications Act*. In doing so, the Commission should distinguish between content and application-neutral traffic management – which may be permissible in some cases – and unlawful application-specific traffic management, which discourages investment in broadband networks, diminishes consumer choice, interferes with users' freedom of expression, and inhibits innovation.

3. The Open Internet Coalition ("OIC") proposes to the Commission this nuanced regulatory approach.[1] OIC represents consumers, grassroots organizations, and businesses working in pursuit of a shared goal: keeping the Internet fast, open, and accessible to all. A list of OIC supporters is attached as Appendix A. OIC is participating in this proceeding because of the key role an open Internet plays, and will increasingly play, for Canadians.

4. In this submission, OIC focuses on the requirement of an open Internet for application developers and their users. It is only because of the open Internet that application developers invest in and can bring the full range of new products and services – the benefits of "innovation without permission" – to users. Similarly, many parties to the Commission's new media proceeding (Notice of Public Hearing 2008-11) noted that an open Internet is the underlying requirement for Canadian content to flourish on the Internet.

5. As summarized by Professor Hogendorn (see Appendix B): "The Internet's incredible societal value derives in part from the fact that it is a general purpose technology (GPT). Rather

---

[1] The Open Internet Coalition's website can be found at www.openinternetcoalition.org.

than being designed only for some limited uses, the Internet is an open platform that acts as an input into a multiplicity of commercial and non-commercial activities, both ordinary and unexpected. Internet users communicate with one another; they email with friends and family, create blogs about politics, develop educational videos, share health resources, and collaboratively generate expansive resources like craigslist and Wikipedia. They create businesses and in turn jobs, building novel applications and enhancing existing tools. In this way, the network's value comes not from delivering traffic in and of itself, but rather from the many ways that Internet users, content providers, and application innovators put the network to use".

6. Sections 27(2) and 36 are two tools available to the Commission to protect the broad public interest in a robust, capacious, open Internet.

7. Although historically s. 27(2) has been applied by the Commission to regulate conduct that is adverse to a Canadian carrier's competitors, it applies equally to prohibit such conduct as directed against any "person" – including an application provider or user. Therefore, if an ISP unreasonably disadvantages the applications of one application provider or a user, such conduct must be prohibited. This approach is consistent with the policy goal of maintaining a competitively and technologically neutral Internet. Using s.27(2) to prevent application-specific traffic management practices is particularly important because the last-mile market for Internet access is highly concentrated. Consumers have minimal ability to "vote with their wallets," by selecting alternative and non-discriminatory facilities-based providers.

8. Section 36 is about removing from Canadian carriers the unrestricted ability to interfere with the choices of consumers concerning what they say and how they say it. It lies at the heart of the Canadian telecommunications system. Because of the important role the telecommunications system plays in facilitating Canadians communicating between themselves and with others, the importance of s. 36 lies not only in economic concepts (the well known efficiency advantages of a common carrier model), but also its manifestation of the principle of freedom of expression. Section 36 requires the Commission to consider in what limited circumstances the public interest indicates ISPs should be allowed to interfere in users' telecommunications.

9. Reading ss. 27(2) and 36 in their ordinary and grammatical meanings, it is clear that they apply to the types of traffic management practices being applied by Canadian carrier ISPs today. The more difficult question confronting the Commission in this proceeding is deciding on what basis some such practices are unjust/undue/unreasonable under s. 27(2) or should be approved by the Commission under s. 36. OIC proposes that these decisions be made on the basis of the following test:

(a) does the traffic management practice further a pressing and substantial objective;

(b) is the traffic management practice narrowly tailored to address this objective; and

(c) is the traffic management practice the least restrictive means to reach the objective.

10. While addressing specific instances of congestion may be a pressing and substantial objective, content, protocol or application-specific traffic management (collectively in this submission, "Application-Specific Traffic Management") practices are neither narrowly tailored nor the least restrictive means to reach the objective. Therefore, the Commission should prohibit Application-Specific Traffic Management.

11. In addition, OIC contributes in this submission the following information to assist the Commission to craft a traffic management regulatory policy:

(a) the availability, feasibility and utility of a variety of traffic management practices other than the application-specific traffic management;

(b) why the natural effect of removing from ISPs the inappropriate crutch of the Application-Specific Traffic Management practices will likely lead to an increase in network capacity;

(c) why it is critically important that ISPs continue to follow global standards; and

(d) the need for ISPs to provide much more robust disclosure to users and innovators in order to allow them to make informed decisions about where to allocate their

resources, how innovators will design their applications, and how users will use them.

12. As well, attached at Appendix B is a paper by Professor Christiaan Hogendorn, prepared for OIC for this proceeding, titled "The Economics of General Purpose Technologies and the Open Internet".

13. OIC intends to make an oral presentation in this proceeding in order to provide to the Commission the benefit of its members' extensive experiences with the Internet and traffic management.

ANSWERS TO THE COMMISSIONS' QUESTIONS

1. Congestion

a) How has Internet traffic grown in the past three years and what are the predictions for its growth in the future? What has been the impact on Canadian ISP networks?

b) How has average end-user bandwidth consumption changed in the past three years and what are the predictions for future changes in Canada?

c) How should congestion be defined in an ISPs network?

d) Are there applications or services that are more likely to cause congestion, and if so, what are they?

e) What are the relative bandwidth requirements for different types of Internet applications?

14. These questions are difficult for OIC to answer because of the poor state of informational transparency relating to traffic on the Internet. Only ISPs have the comprehensive data required to develop the information sought by the Commission. Although the Commission issued interrogatories to ISPs seeking disclosure of data relating to traffic flows, volumes and congestion, the responses to the Commission's interrogatories were extensively redacted, and the information that has been made available on the public record is inadequate for OIC to bring to bear its expertise and help the Commission to answer these questions.

15. Questions 1(d) and (e), the bandwidth requirements of particular applications and services, and their impact on network congestion, should not be examined because ss. 27(2) and 36 preclude Application-Specific Traffic Management practices. The Commission should focus on total traffic volumes and on where exactly congestion, if any, occurs in the network.

2. Technical and Economic Solutions for Internet Traffic Management

a) What technologies could be employed by ISPs (for example, Deep Packet Inspection) to manage Internet traffic?

16. The Commission should decline to accept the false choice between introducing traffic management practices that are inconsistent with the open nature of the Internet and allowing the Internet to become so congested that it ceases to work properly for the majority of users. Canadians can have an Internet that is both uncongested and open, but a nuanced regulatory approach is required in order to achieve that goal.

(i) Investing in Bigger Pipes is the Best Solution for Congestion

17. Before any discussion of specific techniques for addressing congestion it is important to note that the best solution for a congestion problem – the solution that has consistently worked as the Internet has grown – is investing in faster, better networks. Next generation broadband networks not only solve problems of congestion, but are also crucial to promoting innovation and Canadian competitiveness.

18. Limiting ISPs' use of Application-Specific Traffic Management practices (such as throttling certain applications' P2P traffic) will have the beneficial effect of encouraging them to build a bigger and more robust network. OIC is not suggesting that the Commission regulate ISPs to require increased network capacity. However, the natural effect of removing from ISPs the inappropriate crutch of Application-Specific Traffic Management will likely lead to an increase in network capacity. Adding capacity to the network is an important public policy goal. Allowing traffic management practices that encourage ISPs to maintain scarcity in their networks and/or continue to under invest in broadband infrastructure would harm Canada's competitiveness overall and would be contrary to the goals of Canadian telecommunications regulation (s.7(a), (b), (c) and (g)).

19. The most technologically and economically efficient means of managing Internet traffic is by increasing capacity. The advanced networking consortium Internet2[2] confirmed this proposition recently, when it contrasted the introduction of Quality of Service ("QoS") electronics with increasing capacity as a means of addressing congestion.[3] QoS electronics are the hardware that makes Application-Specific Traffic Management possible.

20. Internet2 found that increasing bandwidth is far superior to adding QoS electronics:

> [Increased bandwidth] avoided practical deployment obstacles to implementing any effective QoS across a multiple network environment such as the Internet. Specific obstacles include: coordinating upgrades to QoS technology across every network; changing dramatically network operations, peering arrangements, and business models; and developing suitable means to verify QoS service delivery by users, providers, or both."[4]

21. Internet2 found that the "over provisioning" of bandwidth approach to ensure network performance has been made possible by new technology that provided geometric increases in networking capacity at rates that matched or exceeded Moore's Law.[5]

22. Internet2's experience led it to conclude that increasing capacity is the most economically and technologically efficient means of addressing congestion:

---

[2] Internet2 is a not-for-profit advanced networking consortium comprising more than 200 U.S. universities in cooperation with 70 leading corporations, 45 government agencies, laboratories and other institutions of higher learning as well as over 50 international partner organizations. See http://www.internet2.edu/about/

[3] Beginning in 1998 through 2001, technical leaders from Internet2 worked to develop and deploy an advanced Internet Protocol serviced based on Quality of Service (QoS) technology. This project launched when a large portion of the Internet2 technical community believed that implementing QoS would be a essential to addressing network congestion due to increasing demand for limited bandwidth, especially applications such as streaming video or videoconferencing, which applications do not tolerate packet loss or jitter.

[4] Corbato and Teitelbaum, "Internet2 and Quality of Service: Research, Experience, and Conclusions," May 2006, p.2. See also, Bhagat, Smriti "QoS: Solution Waiting for a Problem". Professor Bhatat's paper concludes that overprovisioning of bandwidth is preferable to QoS technology in addressing network congestion. Available at: http://www.cs.rutgers.edu/~rmartin/teaching/spring06/cs553/papers/004.pdf

[5] Moore's Law refers to the observation in 1965 by Gordon E. Moore, co-founder of Intel, that the complexity of integrated circuits doubles every 24 months with improvements in manufacturing methods.

> Instead of implementing QoS, simply increasing network speed leverages the decreasing cost-per-bit trend of new networking technologies and avoids the pitfalls of QoS implementation. The elegant simplicity of the best-effort service model provided by IP is one of the essential reasons for the success of the Internet. Together with the inherent strengths of connectionless networking and the IP's end-to-end design principle, the best-effort service model has enabled a fast, dumb, cheap, and wildly scalable Internet which has, in turn, provided a foundation for manifold innovative uses, unconstrained by a centralized view of how the network can or should be used.[6]

23. Indeed, data provided by TELUS shows crucial Internet traffic essentially doubled from January 2006 to January 2008.[7] Extrapolating that this growth trend will continue for the foreseeable future, and applying Moore's Law, Internet2's study demonstrates that ISPs should be able to handle such growth without the introduction of QoS electronics (or their related Application Specific Traffic Management practices) as bandwidth capacities will be able to at least correspondingly double over the same period of time.

(ii) Traffic Management Technologies

24. Technologically, there are a number of manners in which ISPs target for differential treatment specific protocols, applications, or content. The Commission should avoid over-analyzing the particular technology used by an ISP but rather focus on the results of the use of any network management tool.

25. Congestion management techniques can be roughly divided between those that are specific to certain protocols, applications[8] or content and those that are not specifically targeted based on any of those criteria.

---

[6] Corbato and Teitelbaum, p. 4.

[7] TELUS(CRTC)4Dec08-1. The TELUS data indicates that the total amount of Internet traffic into and out of the TELUS core backbone network essentially doubled from January 2006 to January 2008. The total megabits per second increased during this time period from 32,390 to 70,651.

[8] Applications encompass computer software and hardware that provide to users processing or communications functionality in addition to that provided by the network itself. Applications provide the intelligence that lies at the "edges" of the Internet and are the reason most users subscribe to Internet access services.

26. As OIC posits in this submission, Application-Specific Traffic Management practices act in breach of ss. 27(2) and 36.

27. One of the technological tools ISPs use to engage in Application-Specific Traffic Management is Deep Packet Inspection ("DPI").[9] DPI involves looking at the content of a communication beyond the header information.[10] DPI devices allow an ISP to inspect the entire content of a communication. This technology also allows the ISP to create, modify, or delete the packets making up a users' communication — and do so at wire speeds — in order to delay, redirect, copy, or block a communication.

28. DPI can be used by ISPs in unacceptable manners. For example, in the United States, Comcast, a large ISP, inserted or "forged" reset packets into their customers' communications, which led to the Federal Communications Commission ("FCC") finding that Comcast's particular use of DPI measures was not "reasonable network management."[11]

29. In contrast, Application-Neutral Traffic Management practices use modern network equipment to allow ISPs to limit the bandwidth use and/or thoughput of their customers, without analysis of the content of the data packets sent and received by the customer and without differentiating among applications.

30. ISPs sometimes employ tools to look at the information contained in the Internet Protocol packet header (the outside of the proverbial "envelope" of a user's communication). This type of monitoring is expected and necessary for the purposes of forwarding a packet to the next hop toward its destination. Such inspection take place at Layers 2 and 3 of the OSI model,[12] which

---

[9] In OIC's view, DPI includes techniques sometimes called Deep Flow Inspection.

[10] Deep Packet Inspection devices have the ability of looking at Layer 2 through Layer 7 of the OSI Seven Layer Model. Bell, Barrett and Cogeco have disclosed to the Commission their use of DPI; see Companies(CRTC)4Dec08-8, BXI(CRTC)4Dec08-8, and Cogeco(CRTC)4Dec08-8.

[11] *In the Matters of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications Broadband Industry Practices and Petition of Free Press et al, for Declaration Ruling that Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for "Reasonable Network Management",* File No. EB-08-1H-1518 and WC Docket No. 07-52, FCC 08-183. 23 FCC Red 13028 41, 46, note 217 (rel. August 20, 2008) ("Comcast Order").

[12] Virtually all networks in use today are based in some fashion on the Open Systems Interconnection (OSI) standard. OSI was developed in 1984 by the International Organization for Standardization (ISO), a global federation of national standards organizations representing approximately 130 countries. The core of this standard is the OSI Reference Model, a set of seven layers that define the different stages that data must go through to travel from one device to another over a network. Under the

contain information about the Type-of-Service (TOS) requested by the end-user's chosen applications – used by routers in the network to decide whether any special handling for a particular packet is appropriate.

31. ISPs can also engage in the "shallow" inspection of a user's communication, which includes looking beyond the network's IP header into the transport-layer header. Shallow packet inspection allows an ISP to examine the TCP or UDP header and not the payload. This form of inspection can examine information at Layers 2, 3, and 4 of the OSI model and is commonly used to view port information (for preventing the propagation of Internet worms by blocking or monitoring traffic on those ports). Some ISPs do this proactively to some degree while other ISPs use shallow packet inspection in response to a complaint. So called shallow packet inspection is sometimes used to identify traffic based on its application, source or content for discriminatory treatment.

32. Examples of Application-Neutral Traffic Management practices include:

- At peak times, artificially allocating bandwidth on a per-user basis (*e.g.,* at certain peak times users would only get 1 MBps of download speed rather than their normal 5 MBps)
- Artificially slowing heavy users at times of congestion (regardless of the application they are using to create that traffic)
- Charging users based on amount of bandwidth consumed to ensure consumers are making efficient bandwidth consumption decisions
- Allowing the current congestion-control algorithms to impact the heaviest users in accordance with the normal and intended operations of the transport's or application's protocols.[13]

OSI Reference Model control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. *See* Webopedia Online Dictionary, The 7 layers of the OSI Model, available at http://www.webopedia.com/quick_ref/OSI_Layers.asp.

[13] Since the late 1980s, the IETF has been aligned behind a packet-conservation method of congestion control In brief, when congestion is detected, senders will immediately and exponentially cut back on their rate of sending and, from that new reduced rate, be conservative in increasing their rates and amounts of sending. Each congestion signal causes the sender to cut sending by

33. The availability, feasibility and utility of such Application-Neutral Traffic Management practices is demonstrated by the fact that they are used by many Canadian ISPs. (OIC suggests the Commission ask the ISPs, by interrogatory, to confirm that they can implement these Application-Neutral Traffic Management practices.) In fact, at least four network operators indicated in their responses to the Commission's interrogatories that they do not employ *any* network management tools at present, demonstrating that Application-Specific Network Management practices are not needed to address network congestion.[14] Moreover, in the United States, Comcast has successfully moved to an Application-Neutral Traffic Management system based on per-user limits.[15]

b) What developments are under way with respect to traffic protocol (such as modifications to transmission control protocols) and/or application changes (such as changes to P2P file exchange) which could assist in addressing network congestion?

34. Nothing works well across a congested network, so it is in the best interest of Internet stakeholders to respond appropriately to a network that is showing signs of stress. This is why OIC recommends to the Commission a nuanced approach to regulation. Current and near-future developments that may address network congestion include:

(a) Today's protocols on the Internet currently already exhibit congestion-control behaviors. If they did not, the Internet would be regularly collapsing as traffic levels increase exponentially year after year while network upgrades are far less regular. If a network product were to be released that always sent at top speed regardless of congestion-control signals, that product would fail to work well because no application works well on a congested path. The traditional and most-

---

50%. Each sender sends fewer packets before receiving acknowledgments. Senders will maintain this new rate for a time and, barring any additional signals of congestion, may slowly increase speed.

These precautions have been designed into all popular TCP/IP implementations in use today and over the past 15 years, It is these methods that have helped the Internet scale during bandwidth constraints over the years. See Congestion Avoidance and Control; November, 1988; Van Jacobson, Lawrence Berkeley Laboratory; Michael J. Karels, University of California at Berkeley; available at http://ee.lbl.gov/papers/congavoid.pdf.

[14] MTS Allstream, TELUS, Primus, and SaskTel have stated that they do not employ network management techniques; see MTS Allstream(CRTC)04Dec08-8, TELUS(CRTC)04Dec08-8, Primus(CRTC)04Dec08-8, and SaskTel(CRTC)04Dec08-8.

[15] See Ex Parte Letter from Kathryn A. Zachem, Comcast Corp., to Marlene H. Dortch, Secretary, Federal Communications Commission, WC Docket No. 07-52, File No. ED-08-IH-1518 (Sep. 19, 2008).

used congestion-control algorithm is known as "Additive Increase, Multiplicative Decrease" ("AIMD") behavior. It is designed to quickly slow down the rate of sending across a network path that is dropping or delaying packets. Once a rate is found that does not result in signs of congestion, a sender can slowly increase speed to probe for faster rates of sending that do not create additional congestion.

(b) The Internet Engineering Task Force (IETF) has already deployed a number of solutions available to users and ISPs to mitigate and avoid congestion. One example is DiffServ (RFC 2474 et al), where users' applications can help identify traffic that is speed-sensitive and ISPs can respond, limit by quota, or ignore such instructions. For example, a residential ISP might offer a quota of 180 MB worth of packets marked "EF" (for "Expedited Forwarding") and the user may use them as they see fit. After the quota is exhausted, packets marked EF will be handled using the standard "Best-Effort" handling (the normal neutral Internet behavior toward packets). This leaves users in charge of deciding traffic priority for themselves. While this method has been available for a long time, ISPs have yet to offer this well-proven technique to residential end-users. Once they do, applications are likely to be designed to use the markings appropriately. Another example is the numerous congestion control standards and methods already published by the IETF as standards or best current practices.

(c) Following the Comcast controversy in the United States, the IETF began investigating additional techniques, some for ISPs, some for end-users and their applications, and some for both, that might result in additional elasticity in links that are awaiting upgrades.

(d) Under the auspices of the Techniques for Advanced Network Applications working group, the IETF is considering proposals that use ISP-provided information concerning the least-costly, least-congested route available from or to particular points on its network. This group will also investigate how to use existing technologies such as data caching to reduce the number of far-reaching connections.

c) What are the specific capabilities offered by the technical solutions identified in (a) and (b) above? For example, would these technologies allow for throttling of individual users or groups of users; would they allow for the collection of information about persons and to what extent?

35. The specific capabilities of the available technical solutions are described above.

36. The capabilities of DPI technologies are worthy of special attention. Like other technology embedded in the telecommunications network, DPI technology is capable of allowing an ISP to listen in on all but encrypted communications. However, the key issue is not the capability of this technology, but instead its effect andwhether, to what degree, and in what circumstances it should be allowed to cause that effect.

37. As is the case in the telephony context, absent due process of law (such as CRTC approval under s. 36, a warrant, or court order), ISPs should be prohibited from accessing the content of communications. If such assurance of privacy is not carried over onto the Internet, then the Internet turns from communications conveyance to public-address system.

38. The Internet is a key platform in today's information economy. For the Internet economy to continue to flourish, users, and businesses must be confident that their Internet communications will be as secure and confidential as their telephone calls. The Commission has a key role to play in building such confidence by regulating the use of DPI technology so that users' communications are not disclosed by ISPs without due process of law.

d) With reference to questions (a) and (c) above, how effective would these solutions be in addressing network congestion in the ISP networks?

39. OIC strongly encourages the Commission to analyze holistically the effectiveness of any measures designed to address network congestion. In other words, the effectiveness of any particular solution must not only measure the degree to which the solution minimizes congestion, but also must measure any collateral damage such solution imposes on the larger Internet ecosystem and users who rely upon the open architecture of the Internet.

e) Also with reference to questions (a) and (c) above, what impact could the implementation of technical solutions have on the Internet Engineering Task Force standards upon which the operation of the Internet is based? Could these solutions create interoperability challenges for application developers?

40. As the Internet is a cooperative of private and public networks, it relies on its various network operators to follow the procedures and standards to preserve interoperability across the network. For example, a software developer in Toronto needs to be able to trust that the behaviors seen on the Internet in her locale will be the same as the behaviors seen in Europe, Africa, Australia, South America, and Asia. These technical standards, which include the basic end-to-end, "best efforts" architecture of the Internet, underpin the ability of developers to innovate locally and deploy their inventions globally.[16]

f) Describe the advantages and disadvantages (including end-user impacts) of employing the following practices in order to manage Internet traffic.

41. In OIC's experience, a combination of some or all of these practices, except practice (vi) (application-based throttling), would be effective in dealing with network congestion.

(i) <u>Monthly Bandwidth Limits (bit caps)</u>

42. As long as they are employed transparently, and there is sufficient competition among carriers, monthly bandwidth limitations do not pose a threat to the open Internet.

43. For the most part residential Internet access has always had a bandwidth limit – or a "cap" – which is an artificial limit programmed into the broadband modem that restricts the amount of information passing that point in either direction. For example, a high-speed Internet account may offer a 20 Gigabyte limit on consumption, but a speed of 50 Megabits per second

---

[16] The end-to-end principle is one of the key architecture principles of the Internet. The principle means that communications and communications protocols should occur at the end points of a Internet without interference from anyone in between. *See End-to-End Arguments in System Design,* Saltzer, J. Reed, D. and Clark, D.D., Second International conference on Distributed Computing Systems (April 1981) pages 509-512, ACM Transactions on Computer Systems, 1984, Vol. 2, No. 4, November, pp. 277-288; Lawrence Lessig & Mark A. Lemley, *The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era,* 48 UCLA Law Review 925 (2001).

(currently state of the art). While compelling, it is illustrative to know that a dial-up modem running at a mere 56 Kilobits per second has a consumption capacity of over 30 Gigabytes.

(ii) Excess Bandwidth Usage Charges

44. As long as they are employed transparently, and there is sufficient competition among carriers, excess bandwidth usage charges do not pose a threat to the open Internet.

(iii) Time-Of-Day Usage Pricing

45. As long as it is employed transparently, and there is sufficient competition among carriers, time-of-day usage pricing does not pose a threat to the open Internet.

(iv) Peak Period Throttling

46. As long as it is employed transparently, and there is sufficient competition among carriers, peak period throttling does not pose a threat to the open Internet ecosystem so long as it is implemented in a manner compliant with ss. 27 (2) and 36.

(v) End-User-Based Throttling

47. If by end-user-based throttling, the Commission means throttling users who consume over a certain threshold of bandwidth, such management does not pose a threat to the open Internet as long as it is done transparently and there is sufficient competition among carriers.

(vi) Application-Based Throttling

48. For the reasons set out at in this submission, application-based throttling for purposes of managing congestion poses a substantial threat to the Internet.

(vii) Content Caching

49. Edge-based content caching has been used for many years and is an appropriate way to improve users' ability to access popular content.

(viii) Upgrading Network Capacity

50. As described above at paragraphs 17 to 23, increasing network capacity is the most economically and technologically efficient means of addressing network congestion. It does not create any adverse side effects for the open Internet.

3. GAS Change Notification Requirements

g) Should these requirements be extended to other ISPs providing wholesale Internet services such as the third party Internet access services offered by cable ISPs?

51. Yes. A competitive market for retail ISPs is crucial to protecting the open Internet.

h) Are similar requirements necessary and appropriate in relation to the provision of retail Internet services?

52. Yes. Robust disclosure creates a better-informed consumer marketplace (where choices are available). Poor disclosure results in an inefficient marketplace, where consumers lack the informed ability to make meaningful distinctions between service providers.

53. That said, the consumer is not the only consideration. The Internet is a cooperative of hundreds of private networks all agreeing to interoperate in a compatible manner. Therefore, robust consumer disclosure only goes so far in solving problems because network developers and operators could not catalog the hundreds of variances from agreed-upon Internet standards that network operators might create.

54. Application providers lack the proper tools to design applications that can efficiently interoperate with all-types of bandwidth constraints or the infinite number of possible bandwidth management techniques of ISPs and transit providers. For these reasons, the Internet was formed around agreement to follow the standards created by the IETF, which openly debates network engineering ideas and formally standardizes the decisions. These standards then become the baseline for networking product innovators to use when designing their products.

55. ISPs often do not provide adequate disclosure to consumers or application providers to allow them to make informed decisions about where to allocate their resources and how to design their applications.

56. The fact that there is currently little transparency concerning traffic management issues is illustrated by ISPs' broad terms of service that generally provide them with the flexibility to change their terms without prior notice to their customers or the public. Even when ISPs provide information to the public, such as in the context of this proceeding, it is far less granular and complete than is needed to achieve the kind of transparency needed by consumers and application providers.

i) If so, what kinds of practices, and/or changes to practices, should trigger these requirements and what information and how much notice should be provided to-end-users?

57. ISPs need to give consumers appropriate disclosure to give them the informed ability to make meaningful distinctions among services and providers. ISPs also need to provide appropriate transparency of network protocols to provide application developers with the ability to design applications that can interoperate with bandwidth limitations of the network operators.

58. The CRTC should require all network operators to publicly disclose:

(a) the specific problem requiring ISP interference, manipulation, or management;

(b) any and all limits imposed on or direct changes made to a customer's upstream or downstream traffic, such as blocking traffic, delaying traffic, deprioritizing or prioritizing traffic, reordering traffic, redirecting traffic, discriminating for or against certain traffic, or inserting traffic into the stream;

(c) technical details of the methods used;

(d) exact details of all thresholds, such as time of day or exact levels of congestion or bandwidth consumption, that triggers any network interference, as well as the

effects on the networks as a result of the chosen thresholds, such as the percentage of users affected and the duration that those users are affected;

(e) exact details of thresholds that trigger a cessation of network interference;

(f) whether and to what extent users' activities and communications are monitored; how that information is used and stored, and with whom is it shared;

(g) the type and nature of data collected, such as dates, times, durations, web or other Internet addresses, TCP packet contents or IP headers; etc; and

(h) prior notice to users of any meaningful changes in terms of service that relate to one of the above-referenced matters.

59. The above information should be collected by the Commission on a periodic and ongoing basis. The Commission should make public as much of the data as possible.

4. Subsection 27(2) – Discrimination, Preferences and Disadvantages

j) What, if any, Internet traffic-management practices employed by ISPs would result in unjust discrimination, undue or unreasonable preference or advantage?

60. Subsection 27(2) of the *Telecommunications Act* provides that "No Canadian carrier shall, in relation to the provision of a telecommunications service or the charging of a rate for it, unjustly discriminate or give an undue or unreasonable preference toward any person, including itself, or subject any person to an undue or unreasonable disadvantage". This provision should be applied to prohibit the use of Application-Specific Traffic Management practices by ISPs.[17]

---

[17] Without doubt, internet service is a "telecommunications service" within the meaning of s.2(1).

(ix) S.27(2) Applies to Conduct vis-à-vis Application Providers

61. Although historically, s.27(2) has been applied by the Commission to regulate conduct that unjustly discriminates[18] or gives an undue or unreasonable preference against a Canadian carrier's competitors,[19] including itself, the provision, taken in its ordinary and grammatical meaning,[20] applies equally to prohibit undue or unreasonable preferences and disadvantages to any "person". As noted by the Supreme Court of Canada, the word "person" has a broad meaning.[21]

62. In other words, Internet users and application providers have standing to make a complaint under s. 27(2). Therefore, if an ISP, for example, subjects an application provider and users to an undue or unreasonable disadvantage by preventing a user from using the application in the manner it was designed to operate, such conduct is prohibited.

63. The application of s.27(2) to Application-Specific Traffic Management practices is particularly important because the last-mile market for Internet access is highly concentrated.[22] Consumers have minimal ability to "vote with their wallets," by selecting alternative and non-discriminatory facilities-based providers. The choice of Internet service from non-facilities based providers is not relevant to this analysis because their services are provided in part over the facilities of ISPs that may engage in Application-Specific Traffic Management practices. While

---

[18] Without doubt, intentionally degrading the traffic of a network user is a matter that can be addressed under s.27(2). See Telecom Decision CRTC 1005-28, paras. 475 and 478.

[19] An example of an Application-Specific Traffic Management practice that could unduly or unreasonably confer a preference on a Canadian carrier would be the throttling of Skype's VOIP service. Though it is a software application, Skype enables its users to make free or affordable long-distance and international calls, thereby effectively competing with some aspects of carriers' business, i.e., long-distance voice calling. Canadian carriers have an incentive to discriminate against Skype traffic in favour of their own traffic. Discrimination against Skype-users by a network operator is thus one type of undue preference arising from control of bottleneck facilities that s.27(2) precludes.

[20] As stated in *Barrie Public Utilities v. Canadian Cable Television Association*, [2003] 1 S.C.R. 476 at para. 20, per Gonthier J.: "The starting point for statutory interpretation in Canada is E. A. Driedger's definitive formulation in his *Construction of Statutes* (2nd ed. 1983), at p. 87:

> Today there is only one principle or approach, namely, the words of an Act are to be read in their entire context and in their grammatical and ordinary sense harmoniously with the scheme of the Act, the object of the Act, and the intention of Parliament.

In the case of federal legislation such as the Act in question, this modern approach to statutory interpretation is confirmed by s. 12 of the *Interpretation Act*, R.S.C. 1985, c. I-21, which provides that every enactment "is deemed remedial, and shall be given such fair, large and liberal construction and interpretation as best ensures the attainment of its objects".

[21] *Barrie Public Utilities v. Canadian Cable Television Association*, [2003] 1 S.C.R. 476 at para. 23, per Gonthier J.

[22] Telecom Decision CRTC 2008-17.

a user can avoid a non-facilities based provider that itself uses application-specific traffic management practices, given the small number of underlying Internet access facilities available to a particular user,[23] he or she may not be able to avoid his or her content being subject to Application-Specific Traffic Management practices.

64. Like all its powers, the Commission's s.27(2) authority must be exercised with a view to implementing the Canadian telecommunications policy objectives.[24] These statutory objectives[25] are consistent with requiring competitively and technologically neutral traffic management practices, particularly when viewed through the lens of the Policy Direction:[26]

> 1. In exercising its powers and performing its duties under the Telecommunications Act, the Canadian Radio-television and Telecommunications Commission (the "Commission") shall implement the Canadian telecommunications policy objectives set out in section 7 of that Act, in accordance with the following:
>
> …
>
> (b) the Commission, when relying on regulation, should use measures that satisfy the following criteria, namely, those that
>
> …
>
> (iii) if they are not of an economic nature, to the greatest extent possible, are implemented in a **symmetrical and competitively neutral manner**, and
>
> (iv) if they relate to network interconnection arrangements or regimes for access to networks, buildings, in-building wiring or support structures, **ensure the technological and competitive neutrality of those arrangements or regimes, to the greatest extent possible, to enable competition from new technologies and not to artificially favour either Canadian carriers** or resellers; [emphasis added]

65. A traffic management practice is not neutral if it singles out particular lawful applications or content for discrimination, preference, or disadvantage.

---

[23] For example, a residential user can usually only choose between the facilities of his or local ILEC or cable company.

[24] *Telecommunications Act,* s. 47.

[25] *Telecommunications Act*, s.7.

[26] *Order Issuing a Direction to the CRTC on Implementing the Canadian Telecommunications Policy Objectives*, SOR/2006-355, December 14, 2006, s. 1 (b)(iv).

66. Under a purposive reading of the language of s.27(2), the *Telecommunications Act* as a whole, as well as the Policy Direction, the Commission must act to prohibit unfair discrimination/preference/disadvantage among users based on the lawful application they choose to use. For example, this means that Canadian carrier ISPs may not unjustly throttle the traffic associated with one application in preference to the traffic associated with another application**.** Such throttling is clearly a differential treatment of different persons (users and application providers) that are in substantially similar situations.[27] Any traffic management measure that degrades the intended functionality of one application relative to any other application should be considered a preference or disadvantage because, by its nature, it harms users and application providers as it

(a) may influence the meaning of telecommunication in some circumstances;

(b) inevitably influences the purpose of the telecommunication by interfering in the users' choice of means of telecommunication; and

(c) influences the future choices of users as to what means of telecommunications they will use and, thereby limits the choices available to them.

67. Also, Professor Hogendorn highlights a number of harms caused by traffic management practices that are inconsistent with the open nature of the Internet (Appendix B):

(a) Traffic management can impact the social value of the Internet, by lessening benefits that "spillover" from both commercial and non-commercial uses;

(b) A focus on Application-Specific Traffic Management can lead to inefficient underuse of applications and protocols, reducing the total social benefit of the Internet;

(c) The risk of adopting Internet applications is heightened if some ISPs may interfere with that application, undermining innovation.

---

[27] See Telecom Decision CRTC 89-3 and Telecom Decision CRTC 2006-7 with respect to the definition of discrimination.

(d) By impeding the use of protocols or applications, an ISP would hamper the use of all the downstream services that depend on the blocked input, and make it less likely that developers would include the input in a new innovation (since there would be uncertainty as to whether the input would work for all users) and make developers less likely to use a given input in the first place (if they fear that its usefulness might be arbitrarily degraded in the future);

(e) Online applications (e.g., instant messaging, or P2P file sharing) that allow direct communication generally become valuable when there are more consumers using an application. When a user either leaves or does not join an application's network (or even uses a network less often), there is a loss of what would otherwise be the positive network effect. A traffic management practice that impedes or degrades the users' ability to access a particular application does not only create a loss for those consumers alone – all of the other users of the application will also receive less value because the direct network effect has been reduced;

(f) Any time an ISP prevents or degrades a certain type of traffic, the effective user base of the relevant application goes down. This will have an impact on new product development, hurting both the developers and the value of other users; and

(g) Similarly, traffic management that differentially affects certain types of services, if these services are partially compatible across platforms, could cause negative effects outside of the Internet.

68. Of course, the most difficult question related to the application of s.27(2) is deciding what conduct is unjust, undue, and unreasonable. In order to assist the Commission to determine if a particular traffic management practice is unjust, undue or unreasonable (collective by, "fair"), OIC offers the following test.

(x) OIC's Proposed s.27(2) Test

69. In analyzing whether a specific traffic management practice complies with s.27(2), the Commission should apply the following three-part fairness test:

(a) does the traffic management practice further a pressing and substantial objective;

(b) is the traffic management practice narrowly tailored to address this objective; and

(c) is the traffic management practice the least restrictive means to reach the objective.

70. Additionally, s. 27(4) places the burden on Canadian carrier ISPs to establish that any discrimination does not violate s. 27(2) and that any preference or disadvantage is not undue or unreasonable.[28]

71. This test is a nuanced approach to the application of s.27(2). For example, it may allow for certain anti-spam and security measures to be undertaken by ISPs while prohibiting the throttling of P2P traffic.

(xi) Application of OIC's s.27(2) Test to Application-Specific Traffic Management Practices

72. Under the first part of OIC's proposed test (pressing and substantial objective), addressing network congestion furthers a pressing and substantial objective – to the extent the Commission concludes such congestion exists at certain times of day and in specific conditions. It is important that the Commission carefully specify the extent of the congestion problem, if any, because only the Commission has access to the information required to make this determination. The more exact the identification of the problem, the more surgical the identified objective. This surgical approach to objective identification will avoid Canadian carrier ISPs using overbroad traffic management practices.

---

[28] Section 27(4) provides that: "The burden of establishing before the Commission that any discrimination is not unjust or that any preference or disadvantage is not unreasonable is on the Canadian carrier that discriminates, gives the preference or subjects the person to the disadvantage."

73. Also, merely because an application is popular does not mean that it harms the network. The addition of on-line images in the 1990s (which caused a strain on the Internet) illustrates this point. The recent trends toward carriage of high-definition pictures and high-fidelity audio could, similarly, be causing a strain today at the most constricted points of the network. Both developments are not harms to be treated – they are growing pains to be endured for a short while and will go away in the course of maturity, i.e., as network upgrades meet the existing and future demands of the users.

74. With respect to the second part of OIC's proposed test, Application-Specific Traffic Management practices are not, by their nature, narrowly tailored.

75. Application-Specific Traffic Management practices are both an overbroad and underbroad solution to the problem of network congestion and, therefore, not narrowly tailored. Such practices are overbroad because they do not consider whether a particular user was actually congesting the network at any given time. They are underbroad because they substitute a particular application as a proxy for heavy use rather than addressing congestion and heavy use directly.

76. The practice of throttling, degrading or blocking a particular application is also not narrowly tailored because of the damage it does to innovation. The Internet's end-to-end architecture allows innovators to create new applications that operate over TCP/IP without having to obtain permission from the network operators that own the transport layer. This has allowed even the most humble start-up company to reach users with few barriers to entry and succeed or fail in the marketplace based on the quality of its product rather than on decisions by the ISP about whether to allow the application to run as intended. Many members of OIC started in garages and as tiny start-ups that grew into globally relevant companies thanks to the innovation-friendly design of the open Internet.

77. Finally, application-specific traffic management practices are not the least restrictive means available to ISPs looking to deal with congestion. As discussed above at paragraphs 17 to 23, in addition to adding capacity, there are several application-neutral traffic management practices that ISPs can employ to address network congestion when it occurs.

5. Section 36 - Control and Meaning of Content

k) What, if any, Internet traffic-management practices employed by ISPs would result in controlling the content, or influencing the meaning or purpose of telecommunications?

78. Section 36 is about removing from Canadian carriers the unrestricted ability to interfere with the choices of users concerning what they say and how they say it. It lies at the heart of the Canadian telecommunications system. Because of the important role the telecommunications system plays in facilitating Canadians communicating between themselves and with others, the importance of s. 36 lies not only in economic concepts (the well known efficiency advantages of a common carrier model), but also its manifestation of the principle of freedom of expression. In other words, s. 36 goes directly to reaching the telecommunications policy objective set out in s.7(a): "to facilitate the orderly development throughout Canada of a telecommunications system that serves to safeguard, enrich and strengthen the social and economic fabric of Canada and its regions".

79. Taking s.36 in its ordinary and grammatical meaning, the language of s.36 makes it clear that it extends past mere content regulation to the means of telecommunication and that the Commission has a statutory duty to vet such practices:

> Except where the Commission approves otherwise, a Canadian carrier shall not **control** the **content** or **influence** the **meaning** or **purpose** of **telecommunications** carried by it for the public. [emphasis added]

80. Parliament has chosen to use in s.36 this variety of words and phrases, which must all be given effect by the Commission. In particular, the provision deals with two related but distinct concepts: (the control of content, and the influencing of the meaning and purpose of telecommunications).

"control the content… of telecommunications"

81. A Canadian carrier may not control the content of telecommunications without Commission approval. Taken individually, the words of this phrase have a wide compass:

(a) Control – (or "regir" in French) means "the power of restraining; the power of directing to exercise restraint or direction over".[29] This verb is clearly not confined to blocking content. However, not all interference with content rises to the level of "control".

(b) Content – (or "contenu" in French) means "the substance or material dealt with (in a speech, work of art, etc.) as distinct from its form or style".[30]

82. Thus, the phrase "control of content" deals with traditional notions of content censorship.

"influence the meaning or purpose of telecommunications"

83. In addition to there being a restriction on the control of content, a Canadian carrier may not influence the meaning or purpose of telecommunications without Commission approval. Thus, Parliament must have intended something more than only traditional restrictions on content censorship. This something more, a prohibition on Canadian carriers having even an effect on the underlying meaning and very purpose of telecommunications, is apparent from the clause chosen by Parliament to express this concept, recognizing that telecommunications are made up of both what a user is trying to express and how he or she is trying to express it.

84. The words of s.36 indicate this Parliamentary intent:

(a) Influence – ("influencer" in French) means "have an effect on".[31] The word "influence" indicates that even effects on the underlying expression and goal of a telecommunication were of concern to Parliament. The threshold for whether a

---

[29] *The Canadian Oxford Dictionary*, 1998 ed., ″control″

[30] *The Canadian Oxford Dictionary*, 1998 ed., ″content″

[31] *The Canadian Oxford Dictionary*, 1998 ed., ″influence″

telecommunication has been the subject of "influence" is lower than that of "control".

(b) Meaning – ("sens" in French) means "what is meant by a word, action, idea, etc; significance; importance".[32] The meaning of a telecommunication is more than its words – it is a more subtle concept than mere content. An audio recording played extremely slowly may retain its content although its meaning is lost because a user cannot understand it. A video streamed with its packets stripped of idenfiers may all arrive on time, but be indecipherable.

(c) Purpose – ("objet" in French) means "a something to be attained; a thing intended; the reason for which something is done or made, or for which it exists".[33] A user's purpose for a telecommunication dictates both what he or she says and how he or she says it. By including the word "purpose" in s. 36, Parliament manifested its intent that all of the factors wrapped into a telecommunication be free from ISP influence and that the focus of s.36 is on the user's judgment concerning the telecommunication. An ISP is never in a position to know a user's purpose in choosing to emit, transmit, or receive certain content by a specific means. For this reason, an ISP is prohibited by s.36 from effecting a telecommunication, including its means of communication, without the Commission approving such an influence on overidding public policy grounds.

(d) Telecommunications – is defined as "the emission, transmission or reception of intelligence by any wire, cable, radio, optical or other electromagnetic system, or by any similar technical system".[34] Thus, the unit of regulation with respect to s.36 is an entire telecommunication and not just the "intelligence" portion. The form of the "emission, transmission or reception" is equally important.

---

[32] *The Canadian Oxford Dictionary*, 1998 ed., ″meaning″

[33] *The Canadian Oxford Dictionary*, 1998 ed., ″purpose″

[34] *Telecommunications Act*, s. 2(1)

85. Section 36 recognizes that what is being communicated and how that content is being communicated are intertwined and both protected from ISP interference. To take a simple example, a live chat message that says "Call the fire department because the house is on fire" would have its meaning (there is currently a fire) and purpose (immediately alert the authorities so that they can put out the fire) influenced if it were delayed by several hours. Clearly it would be unacceptable for an ISP to delay that message. Moreover, only users – and not ISPs – have the legitimacy to decide whether their data is latency sensitive or not.

86. In respect of Application-Specific Traffic Management, users choose applications and protocols that support their specific purposes for the telecommunication of specific content. As discussed above, users, application developers, and content providers expect the Internet to perform in certain ways, based on well-understood global standards and protocols. Influencing the telecommunications associated with specific applications (such as throttling such telecommunications)

(a) may influence the meaning of telecommunication in some circumstances;

(b) inevitably influences the purpose of the telecommunication by interfering in the users' choice of means of telecommunication; and

(c) influences the future choices of users as to what means of telecommunications they will use and, thereby limits the choices available to them.

87. For example, if one evening a user chooses to retrieve from the Internet a specific legal video in order to play that evening, an ISP's "throttling" of that communication influences that telecommunications' purpose since it interferes with its transmission to the extent it is not available to be viewed that evening.

88. Similarly, a user whose purpose is to make available a large collection of legal high-definition, full screen multimedia content may not be able to do so if the innovative application created to allow such a transaction is the subject of Application-Specific Traffic Management. In such a case, both the meaning and purpose of the telecommunications may be influenced by the ISP contrary to s.36.

l) For any Internet traffic management practice identified in (a), what criteria should the Commission apply in determining whether to authorize such practice?

89. At a high level and given the central importance of s. 36 to the Canadian telecommunications system, the Commission should only override the protections of s. 36 when there is a compelling public policy reason to do so.

90. There may be circumstances in which it is appropriate for the Commission to authorize an ISP to control the content or influence the meaning or purpose of a telecommunication. However, that decision should be made based on evidence with respect to the particular practice being proposed by the ISP. The ISP should be required to seek such authorization before it implements the practice at issue.

91. For example, there may be appropriate methods for ISPs to reduce spam in their networks even though doing so may control content and would influence the meaning and purpose of the spam telecommunications. However, by requiring Commission approval before such anti-spam practices are implemented, the Commission would be able to ensure that the proposed solution is appropriate for the spam problem.

92. In determining whether, on the evidence, the Commission should authorize, under s.36, a particular traffic management practice the same three-part test described above at paragraph 69 should be applied. For convenience the test is repeated here:

(a) does the network management practice further a pressing and substantial objective;

(b) is the traffic management practice narrowly tailored to address this objective; and

(c) is the traffic management practice the least restrictive means to reach the objective.

93. Applying that test, for the same reasons as described above at paragraphs 72 to 77, the Commission should conclude that application-specific traffic management practices should be prohibited under s.36.

7. Sections 47 and 7 (The Policy Objectives)

94. As described above, s. 7 and the Policy Direction play an important role in the interpretation and application of the Commission's powers under ss. 27(2) and 36. However, because of the social policy aspects of sections 27(2) and 36, the market forces emphasized by the Policy Direction are not particularly useful.

8. Internet Traffic Management Practices Are a Global Issue

95. In the United States, the FCC has established a Broadband Policy Statement[35] and ruled twice that network operators have unlawfully discriminated against application and content providers.[36] The FCC adopted a policy statement describing principles intended to "encourage broadband deployment and preserve and promote the open and interconnected nature of [the] public Internet."[37] If the Commission allows Canadian ISPs to apply Application-Specific Traffic Management practices to legal content or applications, it would be out of step with US telecommunications policy and would disadvantage Canadian consumers and application providers. Since continued innovation requires a robust, open Internet, Canadians – who would have a second-rate Internet experience – would have less opportunity to develop and deploy innovative Internet applications.

96. The worst outcome of this and similar regulatory proceedings would be for jurisdictions around the globe to permit various Application-Specific Traffic Management practices, thereby

---

[35] *Appropriate Framework for Broadband Access to the Internet over Wireline Facilities*, CC Docket No. 02-33, *Appropriate Regulatory Treatment for Broadband Access to the Internet Over Cable Facilities*, CS Docket No. 02-52, Policy Statement, FCC 05-151 (rel. Sep. 23, 2005).

[36] *Comcast Order, supra*.

[37] The FCC Policy Statement is organized around several principles intended to ensure that consumers are protected from discriminatory or self-interested behaviour by network operators. The Policy Statement, in pertinent part, states:

- consumers are entitled to access the lawful Internet content of their choice

- consumers are entitled to run applications and services of their choice, subject to the needs of law enforcement.

FCC 05-151, August 5, 2005

balkanizing the Internet and threatening the very open, end-to-end architecture that has allowed consumers and application providers to reach across jurisdictions nearly seamlessly.

\5688917

APPENDIX A

List of Coalition Supporters

Adaptive Marketing LLC

Aegon Direct Marketing Services

Amazon

American Association of Law Libraries

American Library Association

Anglebeds.com

Ask.com

Association of Research Libraries

Bloglines

Chemistry.com

Circumedia LLC

Citysearch

CollegeHumor

Computer & Communications Industry Association

Cornerstone Brands, Inc.

Data Foundry

Domania

Downstream

Dreamsleep.com

Dresses.com

Earthlink

eBay

Educause

Electronic Retailing Association

Entertainment Publications

Evite

Free Press

GetSmart

Gifts.com

GoGawGaw

Google

Hawthorne Direct

HomeLoanCenter.com

HSN

IAC

Iceland Health Inc.

iNest

InPulse Response

Internet2

Interval International

iWon

LendingTree

Livemercial

Match.com

Media Access Project

Media Partners Worldwide

Mercury Media

Merrick Group

NationalBlinds.com

Net Coalition

New America Foundation

North Texas Technology Council

PayPal

Product Partners

Pronto.com

Public Knowledge

RealEstate.com

ReserveAmerica

Savvier

ServiceMagic

Shoebuy.com

=

Shopping.com

Skype

Sling Media

Sony Electronics, Inc.

StubHub

Success in the City

TechNet

Ticketmaster

TiVo

Tonystickets.com

Tranquilitymattress.com

US PIRG

Vanguard

Washington Bureau for ISP Advocacy

Windward Instruments

YouTube

APPENDIX B

The Economics of General Purpose Technologies and the Open Internet

Expert Evidence of Professor Christiaan Hogendorn

I. Introduction

The Internet's incredible societal value derives in part from the fact that it is a general purpose technology (GPT). Rather than being designed only for some limited uses, the Internet is an open platform that acts as an input into a multiplicity of commercial and non-commercial activities, both ordinary and unexpected. Internet users communicate with one another; they email with friends and family, create blogs about politics, develop educational videos, share health resources, and collaboratively generate expansive resources like craigslist and Wikipedia. They create businesses and in turn jobs, building novel applications and enhancing existing tools. In this way, the network's value comes not from delivering traffic in and of itself, but rather from the many ways that Internet users, content providers, and application innovators put the network to use.

Because the Internet is a GPT, there are strong economic arguments for sustaining the essential openness of the Internet to maximize its benefits for all Canadians. In turn, there are also strong reasons to be concerned with the bottleneck control that Internet service providers (ISPs) have over how end-users' access the Internet and with network management practices that interfere with or block particular lawful protocols, applications, or content.

The class of GPTs includes the electrical grid, telephone network, and railroads. GPTs are essential to generating economic growth, and the most important uses of GPTs to the economy are often not from the inventor of the GPT itself, but rather by follow-on users who build on top of it. There is a wide body of economic literature that recognizes

the benefits of GPTs, and there are also works that have examined the Internet specifically through the GPT lens.

In this submission, I will provide a snapshot of this literature, in order to outline how it provides a rationale for sustaining the Internet's essential, open character as a flexible input into many activities. I will provide a rough taxonomy of online activities and how they generate particular kinds of benefits to society as a whole.

In so doing, I will highlight how network management practices that interfere with particular protocols, applications, or content -- thus undercutting the general purpose nature of the Internet -- can reduce the Internet's overall value for all Canadians. While an ISP's actions might provide some private benefit for itself, its actions will not properly account for the full impact on society and can dampen the Internet's ability to drive growth.

II. Spillovers

The Internet and GPTs produce substantial social surplus by creating "spillovers": "uncompensated benefits that one person's activity provides to another" (Frischmann and Lemley 2006, pg. 102).

The spillovers from GPTs, and from information technology (IT) in particular, are very large and affect entire economies. Jorgenson and Stiroh (1999) estimate that one sixth of the United States' productivity growth from 1990–96 was attributable to IT. Jorgenson et al. (2008) show that U.S. productivity growth in the early 2000s was based on a wide variety of industries adopting new forms of IT in production. Indeed, extensive research on economic growth and GPTs suggest that economies need GPTs in order to grow (Lipsey, Carlaw, and Bekar, 2005; Jovanovic and Rousseau 2005; Frischmann and Lemley 2006)

Spillovers from the Internet, and from other GPTs, are different from the benefits from special purpose shared facilities like a swimming pool (Frischmann and van Schewick

2007). For these special purpose facilities, the users directly appropriate almost all the benefits, for example, by personally enjoying swimming. In contrast, general purpose facilities create spillovers that impact others. For example, consider a new application for efficient data transfer, which some consumers decides to use for distance learning. The business that created the tool may or may not receive a benefit, in the form of monetary compensation, and the consumers certainly receive a benefit from engaging in this particular activity. In addition, there is a benefit to society as a whole from a more edified population. ISPs benefit in this context, too: by charging consumers a subscription fee, they receive compensation for use of the Internet and capture some value from the consumer. However, the ISP does not capture the full value of consumers' use of this communication tool -- neither the full value in terms of the consumers themselves, and particularly not the spillovers that come from the consumers putting their Internet access to productive uses, and communicating or transacting with someone else.

As I discuss below, the Internet's spillovers fall into two broad categories: innovation and multiple-user spillovers. Innovation includes simple adoption of the Internet (a company becomes more productive by adopting e-mail), creation of new Internet applications (a new website or a new software program), and innovations in the rest of the economy that depend on Internet connectivity as an input (online banking). Multiple-user spillovers ("network effects") occur because so much of the value of the Internet is in different forms of communication and aggregation of its users.

It is important to emphasize here that, while the Internet's value as an input into commercial markets and enhancing productivity is certainly a huge part of its social value, there are many socially valuable uses that are not commercial. The Internet is an input into many public goods as well as nonmarket goods that produce significant spillovers. To consider the impacts of network management practices simply in terms of competition among commercial actors, or the ability of end-users to passively consume content from different commercial actors, misses a whole range of valuable activities, from individuals' blogs, to open source software, to Wikipedia, and beyond.

ISPs benefit from increasing the value of Internet access and enabling many uses of the network, and thus one might ask whether it would be a problem for an ISP to try to capture some of the social surplus for itself, by engaging in business or network management practices that allow it to charge more for particular uses of the network (e.g., use of a particular protocol or application). This might shift the ISP's share of the total benefit of uses of the network, without reducing the total social benefit. Farrell and Weiser (2003) call this *internalizing complementary externalities* or ICE ("complementary externalities" are another term for spillovers).

However, Farrell and Weiser go on to show that, unfortunately, there are a variety of reasons that the self-interest of market actors like ISPs may not align with the public interest, leading them to act in a fashion that is inefficient or anti-competitive. The reasons an ISP might block or interfere with certain services in violation of ICE include (i) the ISP can charge different prices to different customers in order to increase profits at the expense of total market efficiency, (ii) the ISP might try to exclude competitors in the secondary market; (iii) bargaining problems, (iv) the ISP may not fully understand the financial benefits of ICE.

Moreover, Frischmann and van Shewick provide additional reasons why we should not expect an ISP's decisions to interfere with particular applications or protocols to align with maximizing social benefit, particularly in the context of managing network congestion. The spillovers realized through public and nonmarket goods "do not necessarily increase users' willingness to pay for access to the infrastructure resource, and therefore, cannot be appropriated by the network owner through its pricing of the infrastructure good" (Frischmann and van Schewick 2007, pg 402). They also may be diffuse and hard for a network owner to account for. If ISPs ignore these spillovers when interfering with a protocol or application to manage congestion, that will lead to "inefficient underuse" as users will not incorporate the broader social benefits into their decisionmaking (pg. 402). The authors also suggest that non-discriminatory pricing of consumers' use of the network is preferable to "the distortions that result from use restrictions' focus on specific applications and from their all-or-nothing character" (pg 406).

In the next two sections, I discuss the main two types of spillovers generated by the Internet: innovation spillovers and multiple-user spillovers.

III. Innovation Spillovers

The most dramatic source of spillovers from an open Internet is *innovation*. At the broadest level, the Internet is used to increase efficiency and productivity in a wide array of industries. Firms produce new products and find improved ways to make existing products, using the Internet as an input. A bank, for example, can offer financial services with web-based or cell-phone-based updates. A firm like eBay may take a very old activity – running an auction house – and completely redefine it by using the Internet as an input. The Internet can be used for remote medical services and telecommuting, reducing costs. These are examples of extended technological complementarities of the Internet, where the Internet is the key driver, but the spillover takes place in another industry. *Diffusion* like this is common with all GPTs; electricity, for instance, catalyzed significant changes in everything from factory floor layouts to the hours of shopping and working. (David 1990)

Internet applications themselves are a large source of innovation, and they often piggyback on one another. In other words, the Internet is a platform for many innovations, some of which are platforms for innovations as well. For example, consider the World Wide Web itself. It is an innovation built on top of the Internet that in turn has ushered in countless additional technologies -- starting with websites themselves, to search engines, social networking, and so on. Some of these tools might also be platforms for follow-on innovation; for instance, people use Google Maps to create "mash-ups" that combine maps with housing data, and developers create applications that operate on top of Facebook. In this way, innovation online is often "recursive" (Zittrain 2008, pg. 94).

What is the upshot for network management practices that interfere with particular protocols or applications? By impeding the use of a protocol or application, an ISP

would hamper the use of all the downstream services that depend on the blocked input. It would also make it less likely that developers would include the input in a new innovation, since there would be uncertainty as to whether the input would work for all users. Indeed, it might make developers less likely to use a given input in the first place, if they fear that its usefulness might be arbitrarily degraded in the future.

Importantly, the follow-on uses of Internet applications are not always predictable. When Tim Berners-Lee invented the Web, he could never have predicted the wide array of innovations it would give rise to. Thus, whatever the apparent, immediate impact of blocking or throttling a protocol or application today, the full consequences of that cannot be predicted; ISPs' interfering with a particular protocol or application may unwittingly undermine an innovation as far-reaching as the Web.

Another form of innovation may appear less dramatic, but is at least as important to economic growth. Every time a firm or consumer adopts an existing Internet technology, it faces its own unique problems and develops its own unique solutions. For example, a firm may adopt an Internet-based travel expense voucher system for its employees. Obviously many other firms have already done this, but since each firm has slightly different needs, *flexible* technologies are key to a successful adoption. Lipsey, Carlaw, and Bekar note that just because there is a "blueprint" for implementing a technology, this does not include all the "tacit knowledge" that goes with the blueprint. Thus, *adoption* of existing Internet technologies becomes a form of innovation, where each firm has to solve its own problems and implement the technology in its own way. The ability of a firm to accept and implement spillovers from other firms is called *absorptive capacity* (Cohen and Levinthal 1989), and firms that are more open to new technology and more experienced in developing it themselves are generally better at absorbing spillovers.

In absorbing an Internet application, firms face a great deal of risk. There is risk relating to the costs of the project, the benefits of the project, and also the ongoing value of the application in relation to other Internet applications. This relation to other applications is important because most firms try to reuse their applications to perform multiple tasks; this is called *technological convergence* by Rosenberg (1976).

Here, too, discriminatory network management practices can undermine these spillovers. Again, the bare risk of adopting a diffusing Internet application is heightened if some ISPs may interfere with that application. Even if ISPs have not done so yet, just knowing that there is the possibility of interference heightens the risk in absorbing applications. It may also direct the diffusion process away from certain types of applications that would otherwise be preferred.

IV. Multiple User Spillovers

"Network effects" -- an increase in value based on the number of users -- are familiar. For instance, if I am the only one to own a fax machine, then it is not very valuable, but if you buy one, the value of my fax machine increases, and if the whole country owns one, then the value increases further still. Most Internet applications feature these sorts of network effects, creating spillovers.

A. Direct Communication & Network Effects

Many Internet applications allow consumers to communicate directly with each other. E-mail and chat were among the first types of direct communication on the Internet, and today users rely on instant messaging, voice calls, video messaging, and much more. There are also many other Internet applications, often grouped under the heading Web 2.0, that allow users to share profiles and pictures (Facebook), current activities (Twitter), favorite news stories (Digg), and so on. P2P file sharing is also an example of a direct communication between consumers.

All of these types of direct communication become more valuable when there are more consumers using an application. Every additional user increases the number of possible pairs or groups that can communicate, and thus increases the value of the entire system. This means there is a positive spillover generated by each user on the other users.

The reverse of a network effect can be called a *nonuser negative network effect*, so that when a user either leaves or does not join a network (or even uses a network less often), there is a loss of what would otherwise be the positive network effect. (Nagler 2009) Any impediment to a user joining the network therefore has nonuser negative effect on all users. For example, suppose an ISP introduces a traffic management practice that impedes or degrades its users' ability to access a particular application. This does not only create a loss for those consumers. All of the other users of the network will also receive less value because the direct network effect has been reduced.

B. Complementary Applications

In many cases, users of an Internet application do not want to communicate directly with other users, but they do prefer a larger total number of users because this creates a larger market for supporting services or products. This provides one explanation for why avid Apple Macintosh fans try to convince others to switch to Mac – they hope that a larger Mac user base will result in more Mac software. The same logic applies to users of YouTube, online games, or online music services – a larger number of users would lead to more content becoming available.

These spillovers are often called *indirect network effects,* or (more accurately) *complementary components* (Economides 1989). Many economists have studied these component cases (see Shy 2001 for discussion), and generally found that the number of components and social welfare are increasing in the size of the user base, especially if consumers put a high value on product variety.

Any time an ISP prevents or degrades a certain type of traffic, the effective user base of the relevant application goes down. This effect will then impact less new product development, hurting both the developers and the value of other users.

C. Complementarities Across Platforms

Many applications developed for the Internet are also available on related platforms. For example, the financial information service Pageonce is available on both the Web and

the Apple iPhone. Suppose another Internet user joins the web-based service. This will increase the incentive for Pageonce to improve *all* of its product offerings, including the one available on the iPhone, since the protocols used to provide the web-based service are partially compatible with those used to provide the iPhone service. The improved iPhone version of Pageonce generates additional value for iPhone users, so there is a positive spillover for iPhone users that originated with a web-based user. This effect has been called a *cross-platform indirect network effect* since it is a type of indirect network effect that works through the (partial) compatibility of different platforms. (Hogendorn and Yuen, forthcoming)

Because the Internet is at the heart of a constellation of information and communication technologies, these cross-platform indirect network effects can be quite important. They also suggest a further cost to traffic management that differentially affects certain types of services. If these services are partially compatible across platforms, then there could be a nonuser negative effect outside of the Internet on other ICT platforms.

V. Conclusion

The Internet's benefits to society and economic growth are indisputable, and this note has emphasized that most of these benefits come in the form of spillovers – benefits that accrue to third parties in the rest of the economy rather than directly to the ISP or its customer. Spillovers occur beyond the standard buyer-seller transaction represented by a consumer subscribing to an ISP. As with other GPTs, the key to maintaining the value of the spillovers is to maintain the open and *general purpose* character of the Internet connection itself. If ISPs reduce the essentially open nature of the Internet, they can cause considerable harm to society.

About the Author

Christiaan Hogendorn is an Associate Professor of Economics at Wesleyan University, Middletown, CT, USA. His research focuses on market structure and competition in deregulated infrastructure industries. His current projects study network neutrality and other types of "open systems" regulation, and vertical dis-integration in several types of infrastructure industries. He was formerly a Member of Technical Staff at Bell Laboratories and a visiting scholar at the Columbia Institute for Tele-Information (CITI). See http://chogendorn.web.wesleyan.edu/HogendornCV.pdf

References

Cohen, W. and Levinthal, D. (1989). Innovation and learning: The two faces of r & d. The Economic Journal, 99(397):569–596.

David, P. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. The American Economic Review, 80(2):355–361.

Economides, N. (1989). Desirability of compatibility in the absence of network externalities. The American Economic Review, 79(5):1165–1181.

Farrell, J. and Weiser, P. (2003). Modularity, vertical integration, and open access policies. Harvard Journal of Law and Technology, 17(1):85–134.

Frischmann, B. (2005). An Economic Theory of Infrastructure and Commons Management. Minnesota Law Review 89:917-1030.

Frischmann, B. and Lemley, M. A. (2006). Spillovers. Columbia Law Review, 100(2):101–143.

Frischmann, B. and Schewick, B. V. (2007). Network neutrality and the economics of an information superhighway: A reply to professor yoo. Jurimetrics.

Hogendorn, C., Yuen, K., (forthcoming). Platform Competition with 'must-have' components. Journal of Industrial Economics, forthcoming.

Jorgenson, D., Ho, M., Samuels, J., and Stiroh, K. (2008). Industry origins of the american productivity resurgence. Interdisciplinary Information Sciences.

Jorgenson, D. and Stiroh, K. (1999). Information technology and growth. The American Economic Review, 89(2):109–115.

Jovanovic, B. and Rousseau, P. (2005). General purpose technologies. Handbook of Economic Growth, pages 1181–1224.

Katz, M. and Shapiro, C. (1985). Network externalities, competition, and compatibility. The American Economic Review, 75(3):424–440.

Liebowitz, S. and Margolis, S. (1994). Network externality: An uncommon tragedy. The Journal of Economic Perspectives, 8(2):133–150.

Lipsey, R., Carlaw, K., and Bekar, C. (2005). Economic Transformations: General Purpose Technologies and Long-term Economic Growt Oxford University Press.

Nagler, M. (2009). Network externalities, mutuality, and compatibility. papers.ssrn.com.

Papandreou, A. (1998). Externality and institutions. Oxford University Press

Rosenberg, N. (1982). Inside the Black Box: Technology and Economics. Cambridge University Press.

Shy, O. (2001). The Economics of Network Industries. Cambridge University Press.

Trajtenberg, M. and Bresnahan, T. (1995). General purpose technologies:”engines of growth”? Journal of Econometrics.

Zittrain, J. (2008). The Future of the Internet -- And How to Stop It. Yale University Press.