

## Public Comments for CRTC Interrogatory PN 2008-19

---

*Prepared by Christopher Parsons\**

**Summary:** This document aims to identify privacy concerns surrounding the use of Deep Packet Inspection (DPI) devices by Canadian Internet Service Providers (ISPs). After outlining some of these concerns, I note ways that DPI can be used in a minimally invasive fashion, and suggest that content caching might offer a way of alleviating link congestion while avoiding the need to examine data packet payloads.

Version 1.0 :: February 15, 2009

---

\* Doctoral student in the University of Victoria's Political Science department.

## Table of Contents

<b>Introduction .....</b>	<b>1</b>
<b>Overview of Deep Packet Inspection .....</b>	<b>1</b>
<b>Benefits of Deep Packet Inspection for Network Operators .....</b>	<b>2</b>
<b>Privacy Implications of Deep Packet Inspection Devices .....</b>	<b>2</b>
<b>Using DPI Non-Invasively.....</b>	<b>4</b>
<b>Addressing Network Management Challenges .....</b>	<b>5</b>
<b>Conclusion.....</b>	<b>5</b>
<b>Reference.....</b>	<b>7</b>

### Introduction

The Canadian Radio-television and Telecommunication Commission (CRTC) has initiated a public proceeding to consider the Internet traffic managing practices that Canadian Internet Service Providers (ISPs) employ towards retail and wholesale customers. This proceeding is meant to clarify what practices are, and are not, appropriate for managing Canadians' Internet traffic. This author's contribution to the process examines the impact of Deep Packet Inspection (DPI) technologies on Canadians' privacy. After briefly outlining the capacities of DPI technologies, I suggest that their capacity to examine the content of packets poses a privacy risk and, consequently, falls under the CRTC's purview as part of the *Telecommunication's Act*, Section 7(i). I follow by suggesting how DPI could be deployed and avoid infringing/minimally infringe on Canadians' privacy, and conclude by suggesting an alternate method of alleviating network congestion.

### Overview of Deep Packet Inspection

DPI technologies enable network operators, such as ISPs, to examine the payloads of data packets that their customers transmit to, and receive from, the Internet. Using these devices, ISPs can "look inside all traffic from a specific IP address, pick out the HTTP traffic, then drill even further down to capture traffic headed to and from Gmail, and can then reassemble e-mails as they are typed out by the user" (Anderson 2007). DPI devices are designed, in part, to defeat obfuscation techniques some applications use to mask packet contents by drilling past header information and into the packet payload.

Even when consumers encrypt their data traffic, DPI devices can identify the applications responsible for sending and receiving data packets by analyzing how packets are being transmitted to and from a user's computer. By evaluating the spikes and bursts of encrypted web traffic, as well as ports used, it is possible to

correlate traffic patterns with those that particular programs engage in when exchanging data. As a result of this, it is possible to identify traffic, even when the application generating the packets (e.g. Skype, Bit Torrent) actively attempts to subvert packet surveillance (Bonfiglio, Mellia, Meo, et al. 2008).

### **Benefits of Deep Packet Inspection for Network Operators**

By examining packet flows at detailed levels, ISPs can improve network security and guarantee different levels of service to different customer-types. Network security is improved as system administrators can correlate particular packet exchanges with worm- and virus-like behavior, and implement measures to automatically quarantine infected devices from the rest of the ISP's network. Different levels of service can be guaranteed by associating particular application-types with particular usage-plans or priority levels; a user on a VoIP plan might have their VoIP packets prioritized, whereas a user with a Bit Torrent plan might have their packets given priority on the network. Alternately, all individuals may be given the same plan, and simply have some packets prioritized and other deprioritized. I am unaware of any 'application-type' service plans in Canada at the moment, though DPI devices *are* being used to prioritize or de-prioritize packets based on the application that is found generating them. From the ISP submissions for Public Notice 2008-19, one can generally say that DPI-enabled ISPs *are* prioritizing latency sensitive applications, such as VoIP, and de-prioritizing what they have identified as 'time insensitive' applications, such as P2P applications.

### **Privacy Implications of Deep Packet Inspection Devices**

Pursuant to the *Telecommunications Act*, Section 7 (i), it is important that the Commission evaluate whether DPI devices will "contribute to the protection of the privacy of persons." Privacy, when understood as a state that is free from external obtrusions or disturbances, conjoins a series of interrelated though distinctive privacy classifications: freedom to control one's personal information (informational privacy); freedom to isolate oneself (accessibility privacy); and freedom to speak and associate with others (expressive privacy). As a value, privacy is often 'sacrificed' to other interests – interests of security, of profit, of convenience (Bennett 2008, Solove 2008, Torpey 2000) – but in the present proceeding before the CRTC such a sacrifice would have deep impacts on the lives of Canadians. To constrain this discussion, I will focus on how DPI threatens Canadians' expressive privacy.

Judith Wagner DeCew, a noted privacy and legal theorist, has noted that even "surveillance of normal, everyday activities can lead one to be distracted and feel inhibited" (Wagner DeCew 1997: 76). Julie Cohen argues that "[P]ervasive monitoring of every move or false start will, at the margin, incline choices toward the bland and mainstream." Broad-based surveillance, such as the ubiquitous examination of packet streams by ISPs, thus "threatens not only to chill the expression of eccentric individuality, but also, gradually, to dampen the force of our aspirations to it" (Cohen 2000: 1426). This view is shared by Paul Schwartz (2000), psychoanalysts Donald Winnicott (1965) and R.D. Laing (1967), as well as in more

contemporary work, such as that by Daniel Solove (2008) and David Lyon (2008). Further, security experts Susan Landau and Whitfield Diffie have argued that the introducing broad-based surveillance technologies into contemporary packet-based network architectures fundamentally endangers citizens' privacy and raises the specter of vast security risks (Diffie and Landau 2007, Landau 2006).

ISPs who admit to using DPI equipment for parsing data packets are using equipment that has the capacity to examine the payload of packets, or the portions of packets that contain content information of communications. Thus, network operators can examine "Layer-7" of packets, which lets them identify the actual messages that are sent by programs such as Internet Explorer, MSN, or Microsoft Outlook. ISPs who are involved in the CRTC's proceeding have stated that they do not examine the actual content of packets – they have no commercial interest in this – though Bell does acknowledge that they examine the Application layer (Layer-7) of packets. Thus, we can conclude that with the exception of Shaw (who maintain that they are using Arbor-Ellacoya equipment to only examine packet headers), that most Canadian ISPs using DPI devices are already examining packets' Layer-7. They are already reading the 'content level' of data transmissions.

While surveying the application-type that consumers are using to communicate (as has been done by various ISPs, and demonstrated in their ability to classify data-traffic streams in their submissions to the CRTC) arguably does not constitute a violation of individuals' privacy, consumers would not know if by intent or accident an ISP were to record its customers' data-content. Moreover, save for encrypting all of their data traffic, consumers must communicate with the possibility of their ISP engaging in ubiquitous surveillance hanging over their heads. Digital communications are increasingly used by Canadians, and are replacing analogue technologies that they historically have used; email is replacing written letters, instant messages and VoIP replace telephone conversations, and uploading videos to YouTube rather than inviting friends over for a vacation slide show is increasingly normal. These historical, analogue, technologies enjoyed relatively robust privacy protections under Canadian law, and Canadians' digital communications should be similarly safeguarded – technologies that would invade the privacy provided by analogue devices and the laws governing them should be critically examined. Any suggestion that all mail should be examined for content-type by the post office before it is delivered, or that all phone conversation traffic should be persistently monitored though not recorded in the name of managing network congestion would be met with scorn, at best. A similar response should be translated into the digital communicative sphere. Mail systems and telephone networks alike have had to address congestion challenges, but this has not meant that postal outlets or telecommunication providers have inspected the contents of messages and prioritized particular communications-types over others. To have done otherwise would have invaded the privacy of individuals who were communicating with one another. It would have also called into question telecommunications providers' status as common carriers that did not examine content, or discriminate based on its content-type.

It should be noted that none of the publicly filed comments from Rogers mentioned their use of DPI devices to modify their consumers' data traffic. They have used DPI to 'splice' messages into users' data-streams as a technique for managing network congestion; when users approached their monthly bandwidth allotment a message was inserted above web pages customers browsed to and inform them that they were reaching the allotted bandwidth (Stirland 2007). Users were unable to actually *prevent* Rogers from inserting these messages, though were given an opportunity to opt-out of such content injections after receiving their first message. This demonstrates how DPI threatens to undermine the privacy and security of users' communications; such network management technologies open the possibility for surreptitious surveillance and/or alteration of content presented to end-users without their ever being aware of the surveillance, or more subtle alterations. Diffie and Landau's comments surrounding the risks engendered with ubiquitous surveillance of Internet communications ring especially true in the case of ISP deployment of DPI, where a security breach could let third-parties survey and alter Canadians' data traffic.

In summary, DPI technologies endanger Canadians' privacy because it is a technology that can examine and/or analyze the entirety of their non-encrypted data sent to and received from the Internet. This is not to suggest that ISPs are presently using the technology to cache Yahoo! messenger messages, or the video and voice data transmitted using VoIP applications, but at least one Canadian ISP has demonstrated a willingness to use DPI to modify the contents of web pages. The shift to analyzing data content (and in some cases modifying it), as well as dynamically altering transmission bandwidth available to particular applications is an radical shift from the stance of ISPs as common carriers who transmit data traffic with little regard to its content. Previously, only in cases of emergency were particular communications (e.g. police, fire) given priority over other communications – what was once exceptional in the sphere of analogue communications threatens to become the norm in the sphere of digital transmissions.

### **Using DPI Non-Invasively**

Shaw is presently using DPI devices to solely examine data packets' header information – as stated, their devices are not collecting any more data than is normally required to deliver packets to their destination. Were all Canadian ISPs to use their equipment similarly, then their use of DPI couldn't be reasonably considered to be invading Canadians' privacy any more than typical data routing equipment does. It is when ISPs actively penetrate the payload, and in particular Layer-7 of packets, that concerns over privacy violations arise.

If this method of packet analysis is unsuitable, then adopting systems where customer data is clearly segregated from data packet streams is the next-best option. This is already a process that occurs in Bell's system, where customers' personal information (e.g. name, age, etc) are divorced from their customer identification number. As outlined by Bell, their DPI equipment does not correlate

personal information with the consumer identification number, and thus cannot discriminate towards particular customers based on personal information. This leaves unresolved the underlying security challenges brought about by DPI and the accompanying privacy risks, while addressing some concerns that ISPs could survey customers' data traffic. Of course, the disclosure of personal information when browsing the Internet may facilitate the identification of a person to a number.

### **Addressing Network Management Challenges**

In question 2 (f), the Commission asks parties to comment on the advantages and disadvantages of employing content caching to manage Internet traffic. In the publicly released documents, no Canadian ISP has made reference to content caching. This practice has been adopted by major search providers, such as Google, to improve user experiences by effectively 'bringing content closer to end users' – if a number of users in a region are requesting the same YouTube clip, then rather than having to access a server on the other side of the world the video clip would be cached more locally to the users (Whitt 2008). Given that ISPs are identifying P2P applications, such as Bit Torrent, as the key source of link congestion they could adopt a model similar to Google, or that of Comcast, to simultaneously alleviate congestion while enhancing customer experiences.

Comcast, a major American ISP, has experienced similar challenges as those faced by Canadian ISPs. In its case, is trialing P4P's iTracker technology, which "can increase P2P download speeds by 80 percent on ISP networks without materially increasing the network load" (Anderson 2008). P4P is meant to localize P2P files, which means that these files are simultaneously delivered to customers more rapidly while reducing peering and transit links load with other networks. Using P4P, or similar, technologies it is possible to reduce the number of 'hops' that customers' packets must make along an ISP's network and thus reduce network congestion. Comcast engineers have enthusiastically noted that P4P uploads "did not appear to materially increase uploading" for the trialed content (Anderson 2008). Caching works to the benefit of all parties in this case; congestion is managed through the localization of content, and download/upload customers experience increased traffic speeds. By localizing files, it is also possible for customers to more rapidly download content, and reach sharing ratios, thus reducing the time that peering applications must remain open. Rather than adopting a system that discriminates against particular content-types, Comcast is testing a method that avoids such discrimination in favor of a win-win solution. In publicly filed documents, there is no mention that Canadian ISPs who are deploying/have deployed DPI equipment have engaged in similar trials to experiment with similar 'win-win' solutions.

### **Conclusion**

Canadians are increasingly turning to digital telecommunications to communicate with one another, to pursue engage in self-education, and engage in intimate relationships. Actions that once could only occur in physical spaces are increasingly becoming virtual, and it is important that in determining appropriate network management techniques and practices that the freshness of the digital does not

overwhelm our common-sense attitudes concerning appropriate levels of surveillance and analysis of what remain deeply personal actions.

## Reference

Anderson, Nate (2007). "Deep Packet Inspection meets 'Net neutrality, CALEA," *ArsTechnica*. Published July 25, 2007. Last accessed October 10, 2008. Accessible at: <http://arstechnica.com/articles/culture/Deep-packet-inspection-meets-net-neutrality.ars>

Anderson, Nate (2008). "Comcastic P4P trial shows 80% speed boost for P2P downloads," *ArsTechnica*. Published November 3, 2008. Last accessed February 14, 2009. Accessible at: <http://arstechnica.com/old/content/2008/11/comcastic-p4p-trial-shows-80-speed-boost-for-p2p-downloads.ars>

Bennett, Colin (2008). *The Privacy Advocates: Resisting the Spread of Surveillance*. Cambridge, Mass.; The MIT Press.

Bonfiglio, Dario, Marco Mellia, Michela Meo, Dario Rossi, and Paolo Tofanelli (2007). "Revealing Skype Traffic: When Randomness Plays With You," *Computer Communications Review*, vol. 37(4), pp. 37-48.

Cohen, Julie (2007). "Examined Lives: Informational Privacy and the Subject as Object," 52 *Stanford Law Review* 1373.

Laing, R.D. *The Politics of Experience*. New York: Ballantine Books.

Lyon, David (2008). *Surveillance Studies: An Overview*. Malden, MA: Polity Press.

Schwartz, Paul M. (2000). "Privacy and Democracy in Cyberspace." Last accessed February 15, 2009. Available at SSRN: <http://ssrn.com/abstract=205449>

Solove, Daniel J. (2008). *Understanding Privacy*. Cambridge, Mass.: Harvard University Press.

Torpey, John (2000). *The Invention of the Passport: Surveillance, Citizenship and the State*. Cambridge, UK: Cambridge University Press.

Wagner DeCew, Judith (1997). *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*. Ithica, New York: Cornell University Press.

Winnicott, Donald (1965). *The Maturation Processes and the Facilitating Environment: Studies in the Theory of Emotional Development*. New York: International Universities Press.

Whitt, Richard (2008). "Net neutrality and the benefits of caching." Google Policy Blog. Posted December 15, 2008. Last accessed February 14, 2008. Available at: <http://googlepublicpolicy.blogspot.com/2008/12/net-neutrality-and-benefits-of-caching.html>