sandvine
Intelligent Broadband Networks

April 30, 2009

Mr. Robert A. Morin
Secretary General
Canadian Radio-television and Telecommunications Commission
Central Building
1 Promenade du Portage
Gatineau, Quebec
K1A 0N2

Re: Telecom Public Notice CRTC 2008-19 - Review of Internet Traffic Management Practices of Internet Service Providers. Reply Comments.

Dear Mr. Morin,

1. Attached are Sandvine's Reply Comments in connection with Telecom Public Notice CRTC 2008-19, Review of Internet Traffic Management Practices of Internet Service Providers.

Yours truly,

Michael Verhoeve
Sandvine Incorporated ULC
Vice-President, General Counsel

Attachments

c.c.: Parties to Telecom Public Notice CRTC 2008-19

END OF DOCUMENT*

1. Sandvine (or the Company) is a Canadian company located in the heart of the Region of Waterloo technology cluster. The Company was established in 2001 and employs over 250 people in Canada. Sandvine has twice been named to the Deloitte Technology Fast 50 list of fastest growing technology companies in Canada: in 2007 Sandvine was the top company and in 2008 Sandvine was ranked seventh. The Company was identified in the National Post as one of Canada's Top 100 Corporate R&D Spenders, based on fiscal 2007 spending. For the last three years Sandvine has been named one of the top 50 "Best Workplaces in Canada" in Canadian Business magazine.

2. Sandvine is focused on protecting and improving subscribers' quality of experience on the Internet. Sandvine's Network Policy Control equipment and solutions help cable, DSL, FTTx, fixed wireless and mobile operators better serve their subscribers and understand network trends; offer new services; mitigate malicious traffic; manage network congestion; and deliver QoS-prioritized multimedia services. Sandvine's technology is used by more than 150 Internet service provider customers in over 60 countries. Together, these customers serve over a hundred million broadband and wireless subscribers.

3. In connection with Telecom Public Notice CRTC 2008-19, *Review of Internet management practices of Internet providers* (the Notice), Sandvine provides this reply (Reply) in response to certain respondents' initial comments. Sandvine has focused its Reply on the following topics:

    - Privacy
    - Application-specific Policies
    - IETF Standards

4. In the Notice, the CRTC's focus is on "traffic management" practices. Based on the questions in the Notice, the CRTC seems most interested in practices that focus on mitigating congestion in service providers' networks. In this Reply, Sandvine refers to these practices as "congestion management". Certain respondents[1] to the Notice have included comments directed at Internet practices or solutions (such as targeted advertising, lawful intercept and copyright enforcement, etc.) outside the scope of what would be considered "congestion management," or any possible broader definition of "traffic management," and therefore the comments are also outside the scope (as Sandvine understands it) of the CRTC's review as described in the Notice. While it is true that some of these other practices or solutions that have been referenced by certain respondents can potentially use some technological approaches similar to the technological approaches used within congestion management technologies, for example deep packet inspection (DPI) capabilities, it is Sandvine's submission that such cases are very rare today. Further, these other types of practices and solutions are also supportable by a myriad of other technologies, not just DPI.

1 e.g., Office of the Privacy Commissioner of Canada

Privacy

Congestion management solutions don't inspect content

5. Some respondents[2] to the Notice erroneously equate the *ability* to inspect Layer 7 traffic and the *ability* to inspect the "payload" of a packet as de facto inspection of "content" ─ and therefore de facto invasion of privacy. In fact, each layer has a header and payload, all the way up through Layer 7 and beyond, and networking equipment has always read Layer 7 "payload" data. For example, mail servers route mail based on the e-mail address, which is located in the Layer 7 payload data. Session Initiation Protocol (SIP) is a signaling protocol widely used for setting up and tearing down multimedia communication sessions such as VoIP. SIP needs to look in the Layer 7 payload data to find both phone numbers involved in a VoIP conversation, then set up the data (voice) flow. Routers/firewalls look at the Layer 7 SIP exchange to extract this flow information to let the data through. If they don't, the voice component is blocked.

6. Sandvine submits that the true "content" of an Internet transmission is represented as the body of your e-mail message; the music or movie you are downloading; the video you are streaming; the words in your VoIP call, etc. As explained in Sandvine's initial comments to the Notice, Sandvine's congestion management solutions, including those that employ DPI, *do not inspect content* as the content is not relevant to a congestion management solution. To be clear, they:

    - Do not read your e-mail;
    - Do not listen to your voice calls;
    - Do not watch the video you are streaming, etc.

DPI is necessary

7. As described in paragraph 8 of Sandvine's initial comments, DPI is necessary for the identification of traffic today because the historically-used "honour-based" port system of application classification no longer works. Essentially, some application developers have either intentionally or unintentionally designed their applications to obfuscate the identity of the application. Today, DPI technology represents the only effective way to accurately identify different types of applications.

Flow-based and signature-based inspection are both necessary

8. The DPI technology used in Sandvine's congestion management solutions employs two primary inspection techniques: (i) behavioural flow-based; and (ii) signature-based. The techniques can be used separately or together – whichever most

2 e.g., Office of the Privacy Commissioner of Canada, Campaign for Democratic Media, Open Internet Coalition

efficiently and effectively achieves the appropriate level of accuracy in application identification.

9. With behavioural flow-based inspection, which Sandvine uses to identify malicious traffic and certain encrypted protocols (e.g. P2P file-sharing, etc.), the packet itself is not inspected – at all. Instead, the behaviours of packets in a flow are analyzed to identify matches with known application behaviours. For example, a denial of service attack can be readily recognizable by the pattern of "hits" on its target server. As no inspection of the actual packets takes place, there is no opportunity to view any content.

10. Behavioural flow-based inspection does have limitations though. First, on its own, it is only sufficiently accurate in identifying classes of applications that readily display identifiable behaviours. Despite claims of certain respondents[3], any system that relies entirely on behavioural flow-based identification of all applications would be highly subject to an unacceptable level of inaccuracy in identification, resulting in unreliable network demographics reporting and inappropriate and inconsistent policy application. Because behavioural flow-based techniques can generally only identify to the level of application class (such as "bulk" – see paragraph 48 of Sandvine's initial comments) rather than down to the individual protocol level (such as HTTP), any policies created based on the technique may be overly broad and affect more users, applications and protocols than absolutely necessary to achieve the congestion management goal.

11. With signature-based inspection, a library of known application "signatures" is compared to a packet to look for matches. By way of example, the diagram below shows the breakdown of a SIP VoIP packet against the Open Systems Interconnection (OSI) Reference Model, which is an abstract description for layered communications and computer network protocol design.

TCP/IP Layer	Addressing	Equipment Reading	SIP Invite
Layer 1 (wire)	Wire type, location	Provisioning tools	
Layer 2 (link	Ethernet address, Ethernet type, 802.1p tags, vlan tags	Ethernet switch	Ethernet address, Ethernet type, 802.1p tags, vlan tags
Layer 3 (internetwork)	IP address, DSCP, protocol	Router	src IP address: caller dst IP Address: SIP Server
Layer 4 (transport)	TCP, UDP ports	Router, SBC, stateless firewall, NAT	src port: X (random) dst port: 5060 (default SIP)
Layer 7 (application)	URL, SMTP address, POP3 mailbox	Router, SBC, stateful firewall, NAT, IDS	INVITE <sip server name> VIA SIP/2.0 ..... Contact: <caller user info IP, PORT>
(content)	Web page, email body	Anti-Virus, content accelerator	Voice Info

12. In most cases for SIP, the caller will contact the SIP server over port 5060. However, a SIP server can, and often is, configured to work on different ports. Thus, to

3 e.g., Dr. Roberts' comments in the submission of the Coalition of Internet Service Providers Inc.

accurately identify this traffic, a signature based on the IETF RFC standard for the SIP protocol is applied to the packet. In this example, and as shown in the fourth column of the diagram, the solution looks for the presence of "INVITE" followed by some server address then "VIA SIP/2.0". By examining this protocol's header, the solution is able to determine whether the flow is SIP. The solution does not look at the flow's contents, i.e., the voice information, as it is not required to make the protocol identification.

13. If there is no match, then the solution immediately forgets the inspected data and compares the next signature definition in its library to the packets being inspected. Contrary to certain respondents'[4] claims, the entire packet is not "scanned", as if browsing through a magazine. Instead, only those locations that hold identifying signature characteristics are inspected and only to the extent necessary to see if there is a match with the signature profile in the library.

14. In either case, the DPI device never records any of the information past the life of the detection, other than the identity of the protocol, and it only uses this information as an input to decide whether it is relevant for the application of a network policy, such as managing congestion. The process is similar to a mail-sorting machine: the address is matched, the decision is made and the address is then forgotten.

15. For both behavioural flow-based and signature-based inspection, once identification has occurred further inspection not only stops, but the attributes examined in the process of arriving at that identification are discarded. For signature-based inspection, identification can typically happen in the first couple of data packets in a stream. More often than not, those first few data packets don't contain data that would typically be considered the "content" of a transmission, such as the text in an e-mail or the voice in a VoIP call, etc. For example, for a SIP-based VoIP call the first two data packets would be part of the "control flow", which is used to establish call permissions and locations, etc., to initiate the call. Data from the actual conversation would only appear in subsequent packets.

Congestion management solutions don't keep personal data

16. Because typical congestion management solutions do not inspect the actual content of users' Internet traffic, they also cannot record, report on, or store such personal information. As explained in paragraph 62 of Sandvine's original comments, the most "personal" information that Sandvine's congestion management solutions record for an Internet account (i.e, not a particular individual, but the IP address attached to an Internet account, which may include access for many individuals) is aggregate volume usage data, by application or protocol. For example, a typical congestion management solution could report the number of bytes of a VoIP protocol sent and/or received by a given Internet account over a fixed period.

4 e.g., Office of the Privacy Commissioner of Canada

Applications of technologies may raise privacy concerns, not technologies themselves

17. As described above, Sandvine submits that the use of DPI-based congestion management solutions do not create a privacy concern in that they do not inspect content for the purposes of traffic classification, nor is any such information stored within such solutions. Despite this fact, certain respondents[5] claim that somehow the mere presence *of DPI-based technology itself* raises privacy issues, and have called for an outright ban on any such technology. Imagine if this approach were applied to other technologies, such as those supporting cameras. Single Lens Reflex (SLR) technology underlies cameras that take photos at family birthday parties. The same technology has been applied for surveillance of individuals and public spaces. One use of the technology raises privacy issues, the other does not. Nobody questions the value or validity of the camera technology. So why question DPI technology? Privacy concerns properly attach to applications or uses of technologies, not to the technologies themselves.

DPI: A long-standing, ubiquitous technology that fuels innovation

18. Banning the use of DPI, would have far-reaching and damaging consequences across the Internet, where the technology is used extensively. The wireless router in your home probably uses DPI to make sure that time-sensitive packets like VoIP or gaming are delivered quickly, while delaying less time-sensitive packets like e-mail. Firewalls, some built right into popular PC operating systems, use DPI to analyze packets for malicious intent like viruses, trojans, and Spam. Libraries, schools and government institutions rely on their firewalls to protect themselves and their users from attacks. Those firewalls use DPI technology. Load balancers and routers, indispensable hardware that distribute traffic on the Internet and private networks, use DPI to identify where a given packet or URL should be routed and what priority it should be given.

19. DPI is also a key part of the innovation in allowing a migration from IPv4 to IPv6[6], allowing a network operator to convert from one to the other using a carrier-grade network-address-translation (NAT) and keeping protocols such as VoIP operational.

Privacy-sensitive solutions are in demand

20. As described above, Sandvine submits that typical congestion management practices (which the Company believes is the subject of the Notice) do not raise personal privacy issues. However, Sandvine recognizes that other Internet solutions that are in high demand from *consumers, governments and society in general* may raise personal

5 e.g., Campaign for Democratic Media
6 See http://en.wikipedia.org/wiki/IPv6

privacy considerations. Examples, raised by certain respondents[7], include lawful intercept, copyright enforcement, and targeted advertising.

21. To be clear, Sandvine and to its knowledge the majority of "traffic management" vendors in its industry – don't offer such solutions today. However to satisfy demand, clearly some companies need to.

Privacy-sensitive solutions are supported by many technologies

22. To continue the earlier analogy, surveillance of individuals or public spaces could be achieved through a SLR-supported still frame camera or through video recorders supported by a variety of technologies. Similarly, solutions like lawful intercept, copyright enforcement and targeted advertising are achieved through a variety of technologies, not just ─ or even predominantly ─ DPI.

23. Targeted advertising provides a good example. This type of solution can enhance the Internet experience for subscribers by presenting them with more relevant advertising information. Typically, targeted advertising solutions monitor private user Internet activities, such as detailed analysis of website visits, to create a more complete user profile for enhanced marketing. The collection and storage of that type of profile information has clear privacy implications for users, and as expected privacy laws would apply to their use.

24. DPI technology can comprise a component of targeted advertising solutions, but it has been *very rarely* used this way. Instead, other technologies have dominated. Google is one of the leaders in targeted advertising, but to Sandvine's knowledge its targeted advertising solutions do not use DPI. According to Google's own Advertising and Privacy notice in connection with its enormously popular Gmail e-mail application, Google reads your mail to make decisions on targeted advertising: "The Gmail filtering system also scans for keywords in users' emails which are then used to match and serve ads. When a user opens an email message, computers scan the text and then instantaneously display relevant information that is matched to the text of the message.[8]"

25. According to the Google Toolbar Privacy Notice, the Web History service available through the popular Google Toolbar, "records information about the web pages you visit and your activity on Google, including your search queries, the results you click on, and the date and time of your searches in order to improve your search experience and display your web activity. Over time, the service may also use additional information about your activity on Google or other information you provide us in order to deliver a more personalized experience." According to the same Privacy Notice, Google's PageRank service also sends Google "the addresses or other

7 e.g., Office of the Privacy Commissioner of Canada
8 http://www.google.com/privacy_ads.html

information about sites when you visit them.[9]" According to Google's Privacy FAQ, Google stores search engine logs data for each user for 18 months prior to anonymizing it[10]. Again, to Sandvine's knowledge, none of these solutions use DPI.

26. Lawful intercept provides another example of how privacy-sensitive solutions can be enabled by a wide variety of technologies. In the United States under the *Communications Assistance for Law Enforcement Act* (CALEA), service providers are required to identify and intercept criminal data traffic under a lawful warrant provided by law enforcement agencies. DPI technology could be used in a solution designed to support the collection of that data, but so too could a home computer "tapped" into the communications of the individual that is the subject of the warrant.

DPI enables adequate consent for privacy-sensitive solutions

27. In many cases, questions around privacy-sensitive Internet solutions will ultimately come down to the ability to secure sufficient user consent. To date, vendors of privacy-sensitive solutions like targeted advertising have struggled with providing reliable mechanisms for managing user consent. The mechanisms, whether designed as opt-in (where the user must proactively consent to being subject to the solution) or opt-out (where the user must proactively demand NOT to be subject to the solution) have typically been cookies-based. Cookies are "small pieces of text, stored by a user's web browser, that contain the user's settings, shopping cart contents, or other data used by websites.[11]"

28. There are significant problems with a cookies-based system. Cookies can be cleared by the user (purposely or inadvertently), which then erases the "opt-in" or "opt-out" permissions related to a privacy-sensitive solution. Also, cookies are associated with a particular computer's Internet browser, not the user's Internet account. So, if a subscriber uses his Internet account from multiple computers the targeted advertising permissions stored in the cookie do not follow the user between computers. Similarly, if the same user has multiple browsers on the same computer (e.g., Internet Explorer and Firefox), the targeted advertising permissions stored in the cookie do not follow the user between browsers.

29. Fortunately, a better solution to the consent problem is available, through a network-level association between the subscriber's account and his permission settings related to the privacy-sensitive solutions. Regardless of the computer he uses to access his Internet account or the browser that he uses on those computers, the permissions follow the user. Only if the user intentionally changes his account-level privacy permissions could a previously opted-out user be opted-in. Such a solution can be implemented through the use of DPI technology.

---

9 http://www.google.com/support/toolbar/bin/static.py?page=privacy.html&hl=&v=
10 http://www.google.com/privacy_faq.html
11 http://en.wikipedia.org/wiki/HTTP_cookie

DPI-supported policies will offer consumers more choices

30. Service providers are just beginning to explore other uses of DPI that can make their service offerings more attractive to consumers in an increasingly competitive Internet access market. High-speed Internet services are largely offered in the form of flat rate, monthly, unlimited plans. Consumers may be interested in other types of service plans that better reflect the unique ways that they use their Internet connections. Such plans would likely necessitate the ability to differentiate between types of traffic and applications, which in turn would necessitate the use DPI technology as well as other network intelligence tools.

31. For example, "light" Internet consumers may be interested in a service package that ties their fees to their actual Internet usage. But would the consumer want to pay for malicious traffic that affected his usage in a month, or visits to the service provider's customer service portal to address service issues? A DPI-supported policy solution can distinguish between traffic that a service provider would characterize as billable and non-billable so that subscribers are charged in accordance with their expectations. Such a solution could also alert subscribers when they hit or approach pre-determined usage thresholds to help them control monthly spending.

32. Other consumers may be interested in a service package that guarantees a high quality of service for certain frequently-used, latency-sensitive applications, like Internet video gaming or VoIP. A DPI-supported policy solution that can distinguish between different types of traffic and applications is necessary to enable this type of service package.

Application-specific Policies

Application-specific congestion management is highly targeted

33. Certain respondents[12] under the Notice incorrectly claim that application-specific policies:

    a) are not narrowly-tailored;
    b) do not take into account whether the network is congested, or whether a particular user is congesting a network at a given point in time;
    c) are "underbroad because they substitute a particular application as a proxy for heavy use rather than addressing congestion and heavy use directly."

34. These statements reflect an incomplete understanding of the broad capabilities of congestion management solutions.

12 e.g. Open Internet Coalition, Skype

35. In response to point "a" (and as already described in paragraph 55 of Sandvine's initial comments) a policy that is targeted at disproportionate users of bandwidth can become more targeted by applying an application-specific policy as well. For example, by their nature, applications like VoIP, online video gaming and others do not contribute meaningfully to network congestion, but because they are time-sensitive applications, their usefulness to the consumer is greatly impacted by any delays in their delivery. Congestion management solutions allow service providers to create a narrowly-targeted policy that affects:

    - *only* disproportionate users;
    - *only* applications that contribute disproportionately to bandwidth consumption; and
    - *only* applications that are not time-sensitive.

36. Such a policy would minimally impact users' quality of experience, while achieving the congestion management goal. Sandvine is focused on maximizing the user's Internet experience.

37. With respect to point "b" (and as described in paragraphs 11, 55, 59 and others in Sandvine's initial comments) congestion management solutions can:

    - detect when congestion is occurring in the network;
    - identify who is contributing disproportionately to that congestion; and
    - apply the appropriate policies (whether subscriber-centric, application-centric, combinations of the two, etc.) only at those times of congestion.

38. In fact, some of Sandvine's customers are using precisely this sort of congestion management policy today.

39. With respect to point "c" (and as already described in paragraphs 22 through 32 of Sandvine's initial comments), certain applications continue to represent a disproportionate share of network bandwidth and as a result are more prone to contribute to network congestion when network demand is high. Similarly (and as described in paragraph 55 of Sandvine's initial comments), certain users consume a disproportionate share of network bandwidth. It is unclear to Sandvine why the respondents believe that focusing on congestion-causing applications is a mere *proxy* for addressing congestion while focusing on congestion-causing "heavy use" by disproportionate users is not. Both application-centric and subscriber-centric approaches, whether used individually or in combination, can be effective in addressing network congestion. No two service providers have networks that are identically architected, so experimentation is required to achieve "optimal" approaches for each individual deployment.

IETF standards

DPI-supported policies use IETF-approved techniques

40. Certain respondents[13] have suggested that the use of DPI technology is inconsistent with IETF standards. DPI is a technology that is used for inspection. There are no IETF standards for inspection of traffic so to say that DPI either complies or does not comply with IETF standards is meaningless.

41. Further, many IETF standards implicitly require the use of DPI, such as RFC 3489, "Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)"[14], and RFC 2766, "Network Address Translation - Protocol Translation (NAT-PT)"[15].

42. One of the DPI-supported congestion management policies that Sandvine has historically offered service providers is "session management" of P2P file-sharing traffic through the use of TCP Reset packets (RST packets) (see paragraph 53 of Sandvine's initial comments). Despite the claims of certain respondents,[16] there are simply no IETF standards on when or how RST packets should be used. It is further claimed that the RST packets used in session management are in some way "forged" because an RST packet is supposed to mean that "the other end of the connection has failed." While original implementations of RST packets were for this purpose, as with much on the Internet, their use has evolved. For example, most webservers use RST packets *today* as a mechanism for tearing down TCP connections because it is much more efficient than a four-way connection teardown[17]. In short, RST packets are broadly used today and for purposes other than communicating that "the other end of the communication has failed."

43. Interestingly, Sandvine is not aware of a single service provider in Canada who uses a congestion management approach that involves session management with the use of RST packets. In fact, Sandvine estimates that less than five of its 150-plus customers, globally, continue to use this technique, which the Company originally developed in 2002. As Sandvine has created new, more sophisticated and flexible techniques for managing network congestion, customers have naturally migrated to those.

44. Today, to enforce traffic priority Sandvine's congestion management solutions almost exclusively use diffserv marking[18] and weighted fair queuing[19], not RST packets. Both approaches are well recognized by IETF.

---

13 e.g., Dr. Reed's comments contained in the submission by the Campaign for Democratic Media
14 See http://www.faqs.org/rfcs/rfc3489.html
15 http://www.faqs.org/rfcs/rfc2766.html
16 e.g., Dr. Reed's comments contained in the submission by the Campaign for Democratic Media
17 Source : http://www.tcpipguide.com/free/t_TCPConnectionTermination-2.htm
18 Diffserv marking is a "…mechanism for classifying, managing network traffic and providing Quality of Service (QoS) guarantees on modern IP networks." See http://en.wikipedia.org/wiki/Differentiated_services.
19 Weighted fair queuing is a data packet scheduling technique allowing different scheduling priorities for data flows. See http://en.wikipedia.org/wiki/Weighted_fair_queuing

45. For example with respect to diffserv marking, IETF's RFC 4594[20], "Configuration Guidelines for DiffServ Service Classes," suggests that there should be different prioritization for different applications depending on their sensitivity to delay, loss and jitter. They suggest the following categories of services: telephony, telephony/video signalling, multimedia conferencing, real time interactive, broadcast video, low latency data, high throughput data, and low priority data. The RFC continues to discuss the sensitivity of each of the application categories to network conditions.

Recently-developed standards-based approaches could help mitigate congestion, but they are not broadly adopted today

46. As Sandvine pointed out in paragraph 63 of its initial comments, there are a variety of new standards-based approaches to mitigating network congestion. Some have some technical merit, such as LEDBAT, which Sandvine's Chief Technology Officer has worked with the IETF to help develop. But even LEDBAT would need to be very broadly adopted to have a meaningful effect on network congestion. Many application developers have no incentive to implement such standards as doing so guarantees that their applications would receive less priority in times of congestion. Perhaps then it should not be surprising that wide-spread adoption of such standards has not occurred. This challenge alone may constrain the effectiveness of these developments in the short-term and very possibly altogether.

*** End of Document ***

20 http://www.ietf.org/rfc/rfc4594.txt