Before the
Federal Communications Commission
Washington, DC 20554

In the Matter of	)	
	)	
Preserving the Open Internet	)	GN Docket No. 09-191
	)	
Broadband Industry Practices	)	WC Docket No. 07-52

To: The Commission

COMMENTS OF
SANDVINE INCORPORATED

sandvine®
Intelligent Broadband Networks

SANDVINE INCORPORATED
408 Albert Street
Waterloo, Ontario
Canada
N2L 3V3

January 14, 2010

Introduction

About Sandvine

1. Sandvine appreciates the opportunity to provide comments in connection with the Federal Communications Commission's October 22, 2009 Notice of Proposed Rulemaking (NPRM), "In the Matter of Preserving the Open Internet, Broadband Industry Practices" (FCC 09-93), GN Docket No. 09-191, WN Docket No. 07-52.

2. Sandvine was established in 2001 and employs over 400 people in Canada, the United States, Israel and in remote offices globally. Sandvine was recently named to the Deloitte Technology Fast 500 list of fastest growing technology companies in North America.

3. Sandvine is the global leader in network policy control solutions. Our network policy control solutions make the Internet better by protecting and improving the Internet experience for subscribers. The solutions comprise network equipment and software that help cable, DSL, FTTx, fixed wireless and mobile operators understand network traffic and trends, mitigate network congestion, protect the quality of experience for sensitive applications, offer subscribers new services, mitigate malicious traffic, and improve customer service.

4. Sandvine's technology is used by more than 180 Internet service provider customers in over 70 countries. Together, Sandvine's customers serve over 80 million fixed line broadband subscribers and more than 200 million mobile subscribers. The Federal Communications Commission (the Commission) may recognize Sandvine as the supplier of Fairshare Traffic Management, the solution for managing network congestion implemented by Comcast at the end 2008.

5. The Commission has requested comment on a wide variety of issues. Sandvine has limited our comments to just those sections of the NPRM where we believe we offer expertise and meaningful data, specifically:

- C. Codifying the Existing Four Internet Principles
    - Application of rules just to service providers
- D. Codifying a Principle of Non-Discrimination
- E. Codifying a Principle of Transparency
- F. 1. Reasonable Network Management
- H. 3. b. Applicability of Principles to Different Broadband Technology Platforms
    - Application of the Internet Principles to Wireless
        - Application of Nondiscrimination with Respect to Access to Content, Applications and Services, Subject to Reasonable Network Management

Executive Summary

6. The Internet is a commons. The behaviour of each participant affects the other. Applications affect other applications, content affects the network, users affect the network, the network affects applications, and so on. In such an environment it is inappropriate to frame rules, as the Commission has suggested, as obligations of just one participant: network providers. Industry bodies are currently working on a network-based solution to make industry participants accountable for the impacts of their own network activity.

7. The Internet exists for the benefit of its end users – subscribers. Any rules that govern the Internet should therefore be framed in terms of the entitlements of subscribers (as were the original four principles of the Commission's "Open Internet" Policy Statement) and should only be struck if they are of benefit to subscribers.

8. The Commission's proposed bright-line rule prohibiting discrimination is inappropriate because discrimination (or differentiation) among types of network traffic has been and continues to be an effective network engineering technique to ensure that applications and content are delivered with a quality of experience that subscribers expect.

9. An unmanaged network is not a neutral network. Applications and users all make unique demands on the network so that in the absence of active management, certain users and applications would win the battle for scarce network resources a disproportionate amount of the time. Differentiation among Internet traffic has occurred for more than a decade and is recognized as a necessity by the leading industry bodies related to the development of the Internet. Recent "network neutrality" developments in Canada and the European Union have also recognized that applications make different demands on the network and should require differentiated treatment. As more applications and content are delivered over the Internet in the future, and as requirements of applications become more exacting, such differentiation will become even more important.

10. Differentiation supports the Commission's goals in the NPRM. By helping each application get the resources it needs, application-specific differentiation of network traffic would help encourage investment and innovation in applications and related content, as it has done to date. Application-specific network management practices also help to protect subscribers' rights. For example, in times of network congestion a subscriber-specific network management policy that (to quote Chairman Genachowski) "ensures that very heavy users do not crowd out everyone else" can be made less intrusive and less restrictive from the subscriber's perspective by layering on an application-specific component that limits any traffic management to just those applications that are creating network congestion.

11. Differentiation is good, but anti-competitive behaviour is bad. Subscribers could benefit from protection against potential anti-competitive discrimination by network providers, even though very few real-world instances of such behaviour have been

demonstrated. Such protection can be afforded within Sandvine's suggested framework, which includes the five rules entitling subscribers to: send and receive lawful content; run lawful applications and services; connect and use lawful devices that do not harm the network; competition among network, application, and content providers; and disclosure of any network management practices that affect these rights – all subject to reasonable network management.

12. Recognizing the central role of the subscriber in the Internet ecosystem, Sandvine submits that "reasonable network management" should be defined in terms of the user experience for the specific network in question, and should be judged on a case-by-case basis to allow for differences between networks and the dynamic nature of Internet traffic. Sandvine's suggested framework requires that a reasonable traffic management policy should possess a legitimate purpose, be narrowly tailored, be as minimally intrusive as reasonably possible and be auditable. Any ex ante or bright-line rules prohibiting specific network management techniques are not suitable in the ever-changing Internet commons.

13. Sandvine submits that if its framework were put in place today, it would be flexible enough to apply to network providers of all access technologies, including mobile service providers.

Whose Internet is it Anyway?

14. As with any piece of critical social infrastructure such as highways or the electrical grid, Sandvine submits that the Internet exists to serve its end users – Internet subscribers. A highway exists to serve traveler's transportation needs; vehicle manufacturers, road construction companies, pavers and sign manufacturers all offer valuable products and services towards that end, but the traveler's needs define the needs of the infrastructure. Similarly, the Internet does not exist for the benefit of network providers, application providers, or content providers. It exists to serve the Internet subscriber's communication, entertainment, information, and other ever-expanding needs. The subscriber's needs define what the Internet needs to be.

15. In connection with the NPRM the Commission has identified a number of goals that implicitly or explicitly recognize this central role of subscribers. The NPRM lists as its goals:

- to encourage investment and innovation;
- to promote competition; and
- to protect the rights of users, including promoting speech and democratic participation.

16. Investments and innovations by network, application and content providers are only beneficial and successful to the ongoing development of the Internet to the extent that they result in products or services that are consumed by subscribers to those services. Competition among network, application and content providers is desirable because it offers Internet subscribers choice in services, more favourable pricing, and economic alternatives.

17. Sandvine submits that the central role of subscribers' needs in determining the development of the Internet has certain implications with respect to the NPRM:
    - Any rules governing the Internet must be framed in terms of subscribers' needs.
    - Only rules that advance subscribers' needs should be adopted.
    - What is "reasonable" in the context of "reasonable network management" must be viewed in terms of the effect on subscribers.

18. Throughout this submission, Sandvine's recommendations flow from the principle that the Internet must continue to be developed with subscribers' needs paramount.

Codifying the Existing Four Internet Principles

Common Rules for Users of a Commons

19. The Internet has become a key part of the economy's foundation, and as such there are a number of different stakeholders with an interest in its behavior, including subscribers (or "users"), application providers, content providers and network providers.

20. A natural condition of co-dependence or symbiosis characterizes the way these groups interrelate. Network providers need satisfied subscribers who are willing and able to pay for network connectivity. Application and content providers need subscribers to use their services in order to generate revenue. Finally, subscribers want good network performance and access to compelling applications and content. The entire ecosystem depends on balance - each pillar supporting the others to remain viable.

21. In short, the Internet is a commons. Perhaps not surprisingly then, the Internet has fallen prey to the economic dilemma referred to as the "tragedy of the commons": when too many owners are endowed with the privilege to use a given shared resource, the resource is prone to overuse and eventual depletion or destruction[1].

22. So what are the interests of each of the players in the Internet commons? Individual subscribers are concerned about maximizing their own personal utility of the broadband service. There is no incentive for the subscriber to moderate their use of the network without some form of feedback via the service plan definition, cost, structure and enforcement. As will be discussed in greater detail later, a recent study by Sandvine[2] demonstrated that the top 1% of users is responsible for approximately 25% of total residential Internet traffic (measured in bytes) and 20% of users are responsible for 80% of that traffic. Generally, subscribers are not concerned about how their consumption or use of the network is affecting other subscribers using the same network, and even if they were concerned, they have traditionally not been able to monitor their use of the network to understand their impact.

23. Content providers enrich the Internet experience, but certain content providers (which may also be application providers) can also be disproportionate users of the Internet commons. For example, according to Sandvine's study a single website – YouTube – represents roughly 5% of all Internet traffic. Facebook, iTunes, Xbox Live and Xbox Live Marketplace all represent over 1% of network traffic. These content providers want high network performance for the subscribers who access their services. They are not necessarily concerned about how their utilization of network resources affects other stakeholders, including other content and application providers, contending for the same shared resources.

24. Broadband applications have been developed for an ever-increasing set of uses, with great variation in the demands that they place on network resources. Massively adopted personal communication tools like email and instant messaging provide a high degree of value, yet they put a very light load on the network. Sandvine's study showed that residential email traffic consumes less than 1% of total bytes, globally. In

---

[1] "The Tragedy of the Commons." Garrett Hardin. 1968.

[2] Sandvine. *2009 Global Broadband Phenomena*. See http://www.sandvine.com/downloads/documents/2009%20Global%20Broadband%20Phenomena%20-%20Full%20Report.pdf

other words, all the emails sent over residential Internet access networks in the world consume fewer bytes than *any one* of YouTube, Facebook, iTunes, or Xbox Live. By contrast bulk file-transfer or file-sharing applications (such as peer-to-peer (P2P) file-sharing, storage and back-up services and news groups) have typically been used by a much smaller portion of subscribers but represent a much higher portion of network traffic. According to Sandvine's study, P2P file-sharing and storage and back-up services are both top-five applications in terms of the amount of network bytes consumed: combined, they represent over 30% of network traffic. These applications take advantage of Transmission Control Protocol (TCP), one of the core Internet protocols, to use 100% of available network capacity when available, without regard to fairness among users or applications.

25. Network providers are interested in attracting and retaining subscribers to their Internet access services. They do so by increasing the performance of their networks, and by offering their own proprietary applications and content that make their access services more "sticky", such as voice services. By being in a position to both control access to network resources and to offer applications and content, certain public interest groups have identified an opportunity for network providers to act anti-competitively with independent application and content providers, for example through blocking or degrading the independent services. Of course, doing so runs the risk of losing subscribers, who are the lifeblood of network providers. As a result, few examples of such behaviour have occurred in the United States.

26. The Commission proposes to codify the four principles of its existing "Open Internet" Policy Statement (FCC 05-115, adopted August 5, 2005) as obligations of network providers, rather than as statements of subscribers' entitlements, as originally drafted. Sandvine has two concerns with this approach. First, as described above, the Internet exists to serve subscribers, so framing rules in terms of subscribers' needs is entirely appropriate. Second, all stakeholders in a commons need to work cooperatively to maintain its health long term, and, conversely, none should be singled out to guarantee its adherence to core principles. Singling out network providers, as the Commission proposes, presents some obvious problems:

- *Sending and receiving lawful content*. Certain content providers, such as ESPN360, charge network providers (not end users) for the right to have their subscribers access the content. To be clear, that means that subscribers can only access ESPN360 if their network provider has an agreement in place with ESPN360. If network providers are obliged to guarantee that their users can send and receive all lawful content, as the proposed rule would require, would every network provider be obliged to strike an agreement with ESPN360? Also, subscribers' access to lawful content can be blocked or impaired, at least temporarily, as a result of surges in subscriber activity, not through any actions of the network provider. Recent events, such as Michael Jackson's death resulted in widespread reports of subscribers' inability to access related content as web servers and networks were overwhelmed[3]. The event didn't take place during typical "peak" times for broadband networks so capacity could likely have otherwise been available. Networks simply cannot be economically built to allow for such unplanned peaks in demand. Would the Commission's proposed rule require network providers to build their networks un-economically and inefficiently to guarantee that subscribers have unrestricted access to content at all times?

- *Running lawful applications and services*. Certain applications can prevent the effective use of other applications. For example, certain bulk file-sharing or file-transfer applications can saturate subscribers' access link (or shared links) capacity, either in the upstream direction or the downstream or both, depending on the particular network access technology (DSL, cable, mobile) and the individual network architecture. For example, in cable networks, upstream capacity has traditionally been very limited, in DSL networks certain downstream links can be more subject to congestion and in mobile networks upstream and downstream bandwidth is at a premium as it is fixed by the physical properties of the underlying radio spectrum. A small number of bulk file-sharing or file-transfer application sessions can render online gaming, voice-over-Internet protocol (VoIP) calls and other latency- and jitter-sensitive applications at the same network access point unusable.

27. Sandvine foresees similar problems in framing the proposed fifth and sixth rules of nondiscrimination and transparency as obligations of service providers alone:

- *Nondiscrimination.* In the NPRM, the Commission states that it understands the term "nondiscriminatory" to mean that "a broadband Internet access service provider may not charge a content, application, or service provider for enhanced or prioritized access to the subscribers of the broadband Internet access service provider." While Sandvine is not aware of any of its 180-plus service provider customers charging stakeholders in this way, the ESPN360 example described above provides an example where the opposite is true: a content provider is charging a network for the right to offer its subscribers access to content. This restriction on access to lawful Internet content to subscribers of select networks must also be of concern.

- *Transparency*. While increased transparency from network providers represents an important step forward in helping subscribers make choices between network providers and in optimizing their Internet experience, other stakeholders can affect the user experience so the goal of transparency should apply universally. For example, it would be valuable information for a subscriber to understand the bandwidth demands of applications. Certain applications can overwhelm a subscriber's network connection at a moment in time, rendering other applications

[3] MSNBC.Com, *Texts and tweets spread news about Jackson; Twitter, Facebook, cell phone companies, Web sites report surge in traffic*. See http://www.msnbc.msn.com/id/31566668/ns/technology_and_science-tech_and_gadgets

unusable. Armed with this understanding, a user may make different choices about which applications to use at any moment.

28. The Internet Engineering Task Force (IETF) is the open standards organization that works to develop and promote Internet standards, in particular those related to TCP/IP and the Internet protocol suite. With the IETF's recent work on Congestion Exposure (ConEx), it is beginning to investigate how to make all stakeholders accountable for their impacts on the Internet commons. ConEx was discussed at the most recent meeting of the IETF, on November 10, 2009 in Hiroshima, Japan[4] and is described this way:

> "Congestion Exposure (ConEx) is a proposed new IETF activity to enable congestion to be exposed along the forwarding path of the Internet. By revealing expected congestion in the IP header of every packet, congestion exposure provides a generic network capability which allows greater freedom over how capacity is shared. Such information could be used for many purposes, including congestion policing, accountability and inter-domain SLAs. It may also open new approaches to QoS and traffic engineering."
>
> "The Internet is, in essence, about pooling resources. The ability to share capacity has been paramount to its success and has traditionally been managed through the voluntary use of TCP congestion control. However, TCP alone is unable to prevent bandwidth intensive applications, such as peer-to-peer or streaming video, from causing enough congestion to severely limit the user-experience of many other end-hosts."
>
> "We believe these problems stem from the lack of a network-layer system for accountability -- among all parties -- for sending traffic which causes congestion. We propose a metric where IP packets carry information about the expected rest-of-path congestion, so that any network node may estimate how much congestion it is likely to cause by forwarding traffic. A network operator can then count the volume of congestion about to be caused by an aggregate of traffic as easily as it can count the volume of bytes entering its network today. Once ISPs can see rest-of-path congestion, they can actively discourage users from causing large volumes of congestion, discourage other networks from allowing their users to cause congestion, and more meaningfully differentiate between the qualities of services offered from potential connectivity partners. Meanwhile end-hosts may be freed from rate restrictions where their traffic causes little congestion."

29. So, the Internet "industry" is recognizing the *commons* nature of the Internet and the shared obligations of stakeholders that go hand-in-glove with that. Framing the Commission rules in a manner that assigns obligations to only one stakeholder (network providers) would create irreconcilable friction between the Internet's ongoing natural development, as guided by its leading industry body, and the regulations that apply to it.

---

[4] IETF. See http://www.ietf.org/mail-archive/web/int-area/current/msg02041.html

Codifying a Principle of Non-Discrimination

Differentiation is Good. Anti-competitive Behaviour is Bad.

30. In crafting any rules from the NPRM, Sandvine submits that the Commission must be careful to distinguish between very different meanings of the term "discrimination." Discrimination was originally a neutral term. One of the definitions given by Merriam-Webster is "the process by which two stimuli differing in some aspect are responded to differently" – in other words synonymous with "differentiation." However, in the context of the Network Neutrality debate "discrimination" has evolved to take on the meaning of "anti-competitive" behaviour. Differentiation is good. Anti-competitive behaviour is bad.

31. From Sandvine's experience with many of the leading network providers in the United States, Sandvine's solutions are deployed in order to "differentiate" between network traffic to enhance the performance of the maximum number of subscribers for the maximum period of time, *recognizing that the performance needs of applications and their content vary* (a view long-held by technical bodies, such as the IETF (with DiffServ marking[5]) and the European Telecommunications Standards Institute (ETSI) through their work on the IMS and TISPAN standards, including 3GPP TS 29.211[6]. Any Commission rules resulting from the NPRM should encourage network providers to continue such differentiation, while prohibiting only anti-competitive behaviour that arbitrarily shows favour to selected applications, content sources or users over others.

32. Additionally, certain services and applications that consumers would value receiving over their Internet connection are currently not feasible absent differentiation that enables a minimum quality of experience. An example of such services would be telepresence, which is beyond the delivery capabilities of current networks but could be feasible with appropriate traffic differentiation. The Commission should encourage the development of these services and innovations as long as it is not done in an anti-competitive manner.

Unmanaged is Not Neutral

33. Sandvine issued its 2009 Global Broadband Phenomena study[7] in October 2009, based on network data gathered during September 2009. The study consisted of analyzing data from more than 20 cable and DSL service providers' networks totaling 24 million subscribers. The networks were distributed across five regions: North America, Europe, Caribbean and Latin America, Asia-Pacific and Africa. To Sandvine's knowledge, it is the most comprehensive and diverse study of its kind ever prepared.

---

[5] IETF RFC 2474. *Definition of the Differentiated Services Field (DS) in the IPv4 and IPv6 Headers*. See http://tools.ietf.org/html/rfc2474

[6] See http://www.3gpp.org/ftp/Specs/html-info/29211.htm

[7] Sandvine Incorporated. See http://www.sandvine.com/downloads/documents/2009%20Global%20Broadband%20Phenomena%20-%20Full%20Report.pdf

34. This report, and similar reports published by Sandvine in previous years, arrives at one inescapable conclusion: *an unmanaged network is not a neutral network*. There is inherent differentiation in:

- The demands placed on the network by different Internet applications (and users' performance expectations from applications); and
- The demands placed on the network by different users.

Application Demands are Differentiated

35. Sandvine's study showed that the top five categories of applications, by share of aggregate (upstream and downstream) bandwidth, throughout the day are:

- Web browsing
- Real-time entertainment (comprised of streaming audio and video, peercasting, place-shifting, Flash video)
- P2P file-sharing
- Storage and back-up services
- Secure tunneling (e.g., virtual private network traffic)

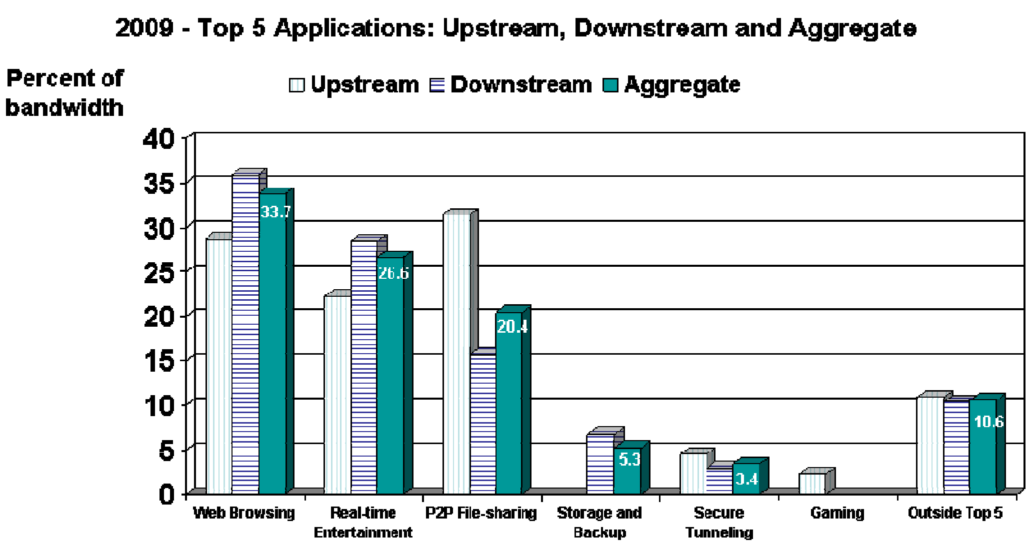

36. A couple of observations can quickly be made. First, users consume these applications in differentiated ways and the applications place differentiated demands on the network. Real-time entertainment is consumed as it is received. For example, streaming video and audio is watched and listened to the instant it arrives at a user's computer. These applications are time-sensitive. Delays in their delivery would be noticed (and not welcomed) by users. P2P file-sharing and storage and back-up services are consumed very differently. Users initiate the process to download or upload files then can walk away from their computers, often overnight. Delays in delivery of the data would not be of particular concern to (or even recognized by) the user.

37. Second, where in the "Top 5" list are email, instant messaging, VoIP and online gaming? Despite the enormous popularity of these applications among Internet users, they don't consume much bandwidth. Text-based data like that in emails and instant messages are simply not bandwidth-hungry, neither are the simple voice data in phone calls or the "move left", "move right" and "fire" commands of an online video game. Yet the performance of these immensely popular applications can be impacted significantly by less popular applications (measured in number of users) that consume much more bandwidth, such as P2P file-sharing or storage and back-up services. Too much delay in a VoIP call and the callers start stepping on each others' words. Much delay in a "fire" command could leave the shooter fired upon: game over.

38. Sandvine's study also demonstrated that during peak hours (shown by the study to be 7:00pm to 10:00 pm, normalized by time zones), when network congestion is most likely to cause performance issues, time-sensitive applications represent an even larger component of network traffic.

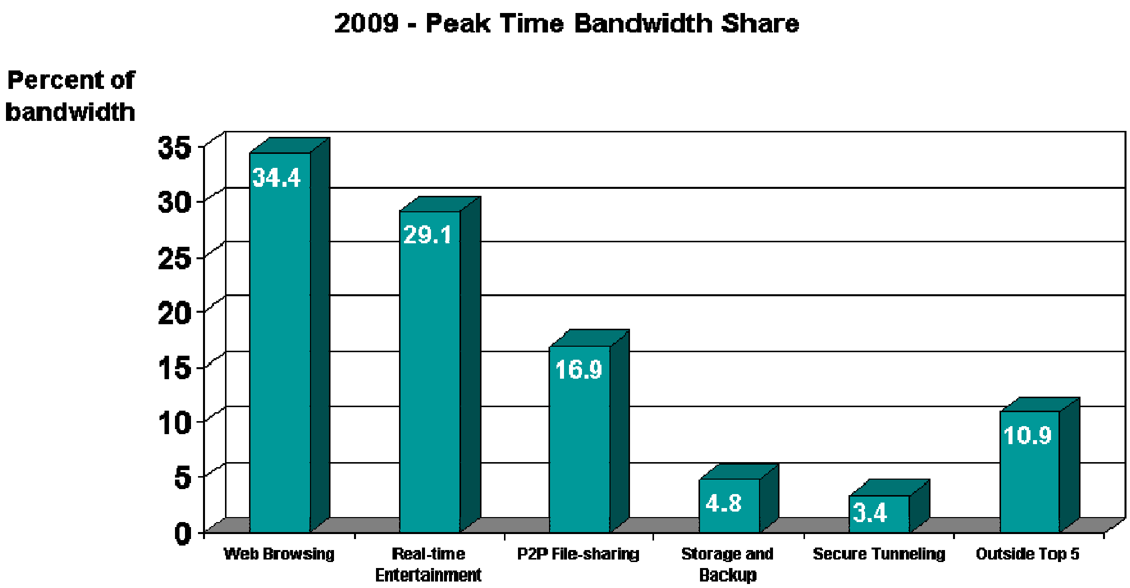

39. The following graph demonstrates how, at peak Internet hours the use of time-sensitive applications surge the most. The peak-time bump in traffic is almost completely attributable to the surging evening popularity of Real-Time Entertainment and Web Browsing – not only do both of these categories experience huge per-subscriber increases in bandwidth demand (rising by almost 35% and 26%, respectively), but these categories also make up a significant portion of the overall utilized bandwidth (29% and 34% of network traffic, respectively).

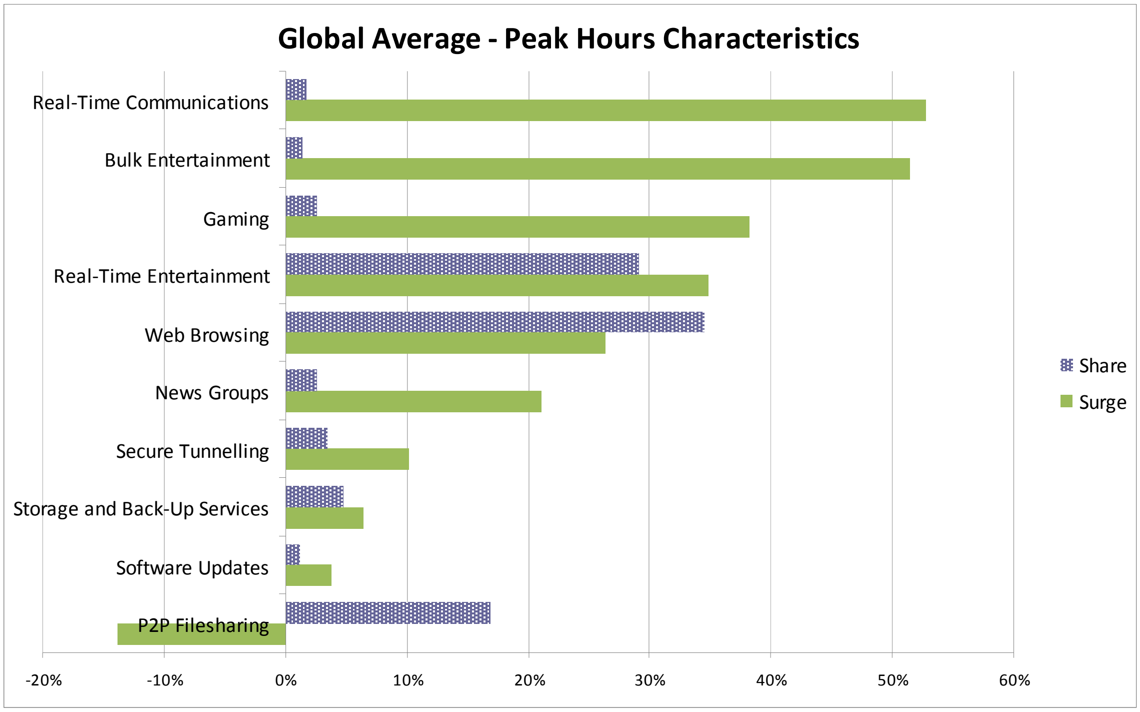

40. The developers of the Internet have long understood the need for differentiated treatment of applications. The IETF's *1998* Request for Comment (RFC) 2475[8] discusses an "Architecture for Differentiated Services," or "DiffServ" as it has come to be known, which explicitly describes methods to provide differentiated service levels to various applications, based on their heterogeneous network requirements. The overview of this RFC describes its purpose as follows:

> "This document defines an architecture for implementing scalable service differentiation in the Internet. A "Service" defines some significant characteristics of packet transmission in one direction across a set of one or more paths within a network. These characteristics may be specified in quantitative or statistical terms of throughput, delay, jitter, and/or loss, or may otherwise be specified in terms of some relative priority of access to network resources. Service differentiation is desired to accommodate heterogeneous application requirements and user expectations, and to permit differentiated pricing of Internet service."

41. Today, DiffServ is the primary protocol routers use to provide different levels of service. Despite this fact, DiffServ marking has a few problems. It is not universally obeyed by current generation access devices in either DSL networks (ATM L2 backbone), or cable networks (DOCSIS layer), where congestion is highest, due to technical limitations in the devices. Also, the devices can't trust the marks: over time applications have cheated the system by mischaracterizing their traffic in order to

[8] IETF RFC 2475. See http://www.ietf.org/rfc/rfc2475.txt

achieve higher priority. This problem created the need for devices capable of deep packet inspection (or "DPI" - a key area of Sandvine expertise) which can reliably identify Internet application traffic.

42. DPI is, from a network engineering and architectural perspective, the act of any network equipment which is not an endpoint of a communication using any field other than the layer 3 destination IP address for any purpose. DPI has been used for over a decade in providing differentiated treatment of network traffic.

43. Home wireless routers use 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.

44. DPI is also a key part of the innovation in allowing a migration from IPv4 to IPv6[9], 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.

45. So, differentiated treatment of applications (through DiffServ marking, DPI and a variety of other means) is not only necessary for the optimal future development of the Internet, but has been a key component in helping the Internet achieve what it has to date.

46. In connection with application-based differentiation of network traffic, the Commission has asked which applications might require a guaranteed quality of service. Applications differ with respect to the amount of bandwidth (or throughput), latency, jitter, and packet loss that they require in order to be delivered at an expected quality of service level. Sandvine submits the following definitions for purposes of the NPRM.

- Bandwidth: traffic volume over time. It is usually measured over a short time, such as bits/second or megabits/second (Mbps), which is 1,000,000 bits/second.

- Latency: the delay for a message to get from one communications end point to the other, e.g., the time it takes for a VoIP data packet to leave the speaker's mouth and arrive at the listener's ear. It is typically measured in milliseconds.

---

9 IETF RFC 2460. *Internet Protocol, Version 6 (IPv6) Specification* . See http://www.ietf.org/rfc/rfc2460.txt

- Jitter: the variation in the latency of one message to another, typically measured in milliseconds (e.g. if the first message takes 1ms and the second message takes 10ms, then there is 9ms of jitter).

- Packet loss: occurs when one or more packets of data traveling across a computer network fail to reach their destination. Packet loss can be caused by a number of factors, including signal degradation over the network medium, oversaturated network links, corrupted packets rejected in-transit, faulty networking hardware, faulty network drivers or normal routing routines.

47. While bandwidth gets most of the attention, adding bandwidth is not always (or even mostly) the answer to improving the user's quality of experience for an application. The other factors can play a crucial role. An application can be classified into one of three categories based on its requirements of a network across these three characteristics:

- Bulk applications. These applications include P2P file-sharing (e.g., BitTorrent, FastTrack, etc), web surfing, usenet news (NNTP), and file transfers over FTP or HTTP, for example, and will go as fast as the network will permit. TCP is designed to achieve the maximum communication rate possible, using all resources available. In practice bulk applications will go as fast as the thinnest part of the network between the client and server. In the case of the server co-located within the ISP network (e.g. a content-delivery network, a cache), this will be bound by the access equipment speed. In the case of a server which is located farther away, this may be bound by transit (connection to all worldwide public networks) or peering (connection to other nearby private networks) performance. Typically, servers of bulk applications (e.g. Speedtest.net, Rapidshare.com, Megaupload.com) will saturate the download speed of the consumer's modem, as they typically download-only. In the case of P2P, it is bi-directional so it can also have the same effect in the upstream direction.

  Most bulk applications can run unattended by the user. File transfers are initiated by the user, who may then walk away – often for hours or even overnight – while the process completes. Bandwidth is the primary determiner of transfer speed of long-running connections and performance will generally improve linearly with increases in bandwidth. As a result, latency and jitter matter much less – users likely would not even notice their effect. Packet loss affects throughput: as packet losses occur, TCP reduces the number of packets sent per second and only increases throughput once packets stop being dropped.

  Web surfing represents an exception in the Bulk category. "Web 2.0" sites have introduced interactive components to web surfing – the user typically attends the activity and data travels bi-directionally as users have started to be content providers in their own right. Increases in bandwidth do not translate linearly to increased performance as latency is a gating factor to the end-user experience. Beyond the time it takes to transfer the content, loading a website typically involves four "round trips" between a personal computer and the related web server. First, the Domain Name Server (DNS) must resolve the domain name (i.e, translate www.sandvine.com to its numeric Internet address) then the three-way handshake established by TCP must be completed[10].

Each of the four round trips is subject to the latency in the network, and when added together this delaying effect becomes the limiting factor in the transmission such that additional bandwidth does not dramatically improve loading times for a website. To illustrate with an example, if the latency in a single round trip is 500 milliseconds (0.5 seconds), a website would take at least two seconds (4 round trips x 0.5ms) to load even if the subscriber had an infinitely high bandwidth connection. Typically "Web 2.0" sites have more than one file (images, videos, ads, etc), so the time to load can be substantially worse.

- Interactive applications. These applications are paced by the consumer. In the case of VoIP, bandwidth largely depends on silence suppression and the codec bandwidth chosen, but it is typically 8-30Kbps. The bandwidth requirements of interactive applications are often modest (though in the case of video conferencing the rates are significantly higher: 200-500Kbps is common), but they typically require very low latency, jitter and packet loss to achieve a satisfactory quality of experience. For example, a VoIP user can perceive latency of 150 milliseconds on a call, and delays greater than 300 milliseconds render the call unusable[11]. As with web surfing, adding bandwidth will not necessarily address quality of service issues. In general, because of the sensitivity of Interactive applications to latency, jitter and packet loss it is particularly important to protect the quality of service for these applications.

- Paced/Burst-paced applications. Streaming video and audio applications such as YouTube and SHOUTcast fall into this category. The media involved has a natural bit rate, and the connection tries to achieve this rate on average over its lifetime, though for short durations the media will 'burst' to provide buffering on the client to allow for packet loss on the network (YouTube, because it uses TCP, will attempt to transmit at line rate initially). So, these applications can be modeled by the media they carry. For typical Internet streaming today, rates of approximately 300-400Kbps are common. Hulu, YouTube, and others are starting to shift to higher-definition video, for which the rate can increase to 1-6Mbps of bandwidth.

  With paced/burst-paced applications it is important that a network sustain the minimum bandwidth requirements, but because of the buffering involved additional bandwidth only marginally improves performance, by making the applications less sensitive to latency, jitter and loss in the network.

---

[10] IETF RFC 793. See http://www.ietf.org/rfc/rfc0793.txt

[11] International Telecommunication Union. *ITU-T Recommendation. G.114*. See http://www1.cs.columbia.edu/~andreaf/new/documents/other/T-REC-G.114-200305.pdf

48. The following table provides some representative benchmarks to achieve a minimum quality of service for certain popular applications. Such figures require significant assumptions, which Sandvine has included as Appendix 1:

Application Category	Application Class	Minimum Bandwidth	Maximum Latency	Maximum Jitter	Maximum Loss
Bulk	P2P	19Kbps	n/a		
	Web surfing	1Mbps (Web 2.0)	166ms (latency + jitter)		n/a
	Email	60Kbps	n/a		
	Usenet news	195Kbps	n/a		
	FTP file transfers	195Kbps	n/a		
Interactive	VoIP	16Kbps	300ms (latency + jitter)		< 0.5%
	Video gaming	50Kbps	75ms (latency + jitter)		< 0.5%
	Video Conferencing	250Kbps	300ms (latency + jitter)		< 0.05%
Paced (and burst-paced)	Video streaming streaming	300Kbps, to not have much of a wait time	< 1s for "channel change"	<50ms	<0.05%
	High def video	1-3Mbps depending on quality of HD.	< 1s for "channel change"	<50ms	<0.05%
	Audio streaming	Audio:160Kbps for CD quality.	< 1s for "channel change"	<50ms	<0.05%

49. These inherent differences in application traffic are starting to be recognized in so-called "network neutrality" decisions around the world. In October 2009, the Canadian Radio-television and Telecommunications Commission (CRTC) concluded its Review of Internet Traffic Management Practices (ITMPs). As part of its decision, the CRTC stated:

> "The Commission notes that the degree to which an application or service is delayed may have an impact on its performance. Furthermore, transmission delays may affect some types of applications or services more than others. For these reasons, it is important to identify which types of traffic and/or applications would be impacted by transmission delays.
>
> In the case of time-sensitive audio or video traffic (i.e. real-time audio or video such as video conferencing and voice over Internet Protocol (VoIP) services), ITMPs that introduce delays or jitter are likely to cause degradation to the service. The Commission considers that when noticeable degradation occurs, it amounts to

controlling the content and influencing the meaning and purpose of the telecommunications in question."

50. Accordingly, the CRTC required that the CRTC vet in advance any network management practice that affects time-sensitive traffic in this way.

51. With respect to non-time-sensitive traffic, the CRTC further decided:

> "With respect to non-time-sensitive traffic, the Commission considers that the use of ITMPs that delay such traffic does not require approval under section 36 of the Act. However, the Commission is of the view that non-time-sensitive traffic may be slowed down to such an extent that it amounts to blocking the content and therefore controlling the content and influencing the meaning and purpose. In such a case, section 36 of the Act would be engaged and prior Commission approval would be required."

52. Similarly, in November 2009 the European Union concluded its Telecoms Reform package, including some new guarantees for an open and more neutral Internet. These stipulations also recognized the differentiated needs of applications or "services" delivered over the Internet:

> "The new telecoms rules will ensure that European consumers have an ever greater choice of competing broadband service providers. Internet service providers have powerful tools at their disposal that allow them to differentiate between the various data transmissions on the internet, such as voice or 'peer-to-peer' communication. Even though traffic management may allow premium high-quality services (such as IPTV) to develop and can help ensure secure communications, the same techniques may also be used to degrade the quality of other services to unacceptably low levels or to strengthen dominant positions on the market. That is why, under the new EU rules, national telecoms authorities will have the powers to set minimum quality levels for network transmission services so as to promote "net neutrality" and "net freedoms" for European citizens."

53. Sandvine has over 180 network provider customers in more than 70 countries. These networks cover more than 80 million fixed line broadband subscribers, which represents approximately 20% of the world's total fixed line broadband subscriber base. Sandvine estimates that approximately 90% of those have deployed application-specific network management policies. Increasingly, there is growing global acceptance that differentiation of network traffic by application is a necessary and effective approach.

Subscriber Demands are Differentiated

54. Just as applications make differentiated demands on the network, so too do users. Sandvine's 2009 Global Broadband Phenomena study demonstrated that over a month the top 1% of users, by total consumption, account for 25% of total bytes on

residential access networks and the top 20% of subscribers account for fully 80% of total bytes. This study also demonstrates that the monthly data consumption of a heavy Internet user exceeds that of an average user by a factor of about 200.

55. Two key factors contribute to this enormous variation in individual network requirements between "top" and "average" subscribers:
    - The network's top users exhibit relatively little change in usage throughout the day. Conversely, average subscribers pop online and offline throughout the day and over the course of a month.
    - Top subscribers still rely heavily on bulk download applications like P2P file-sharing, storage and back-up services, and news groups – applications that are responsible for massive amounts of traffic volume with very little user involvement.

56. The graphs below show the top five categories for a top subscriber and an average subscriber over the course of the day. The two aggregate bandwidth graphs share a common y-axis scale, so visual comparisons between them are valid. Since the scales are consistent, we know that at any instant in time, a top subscriber is likely to be using more bandwidth for P2P file-sharing than an average subscriber uses in total for all categories.

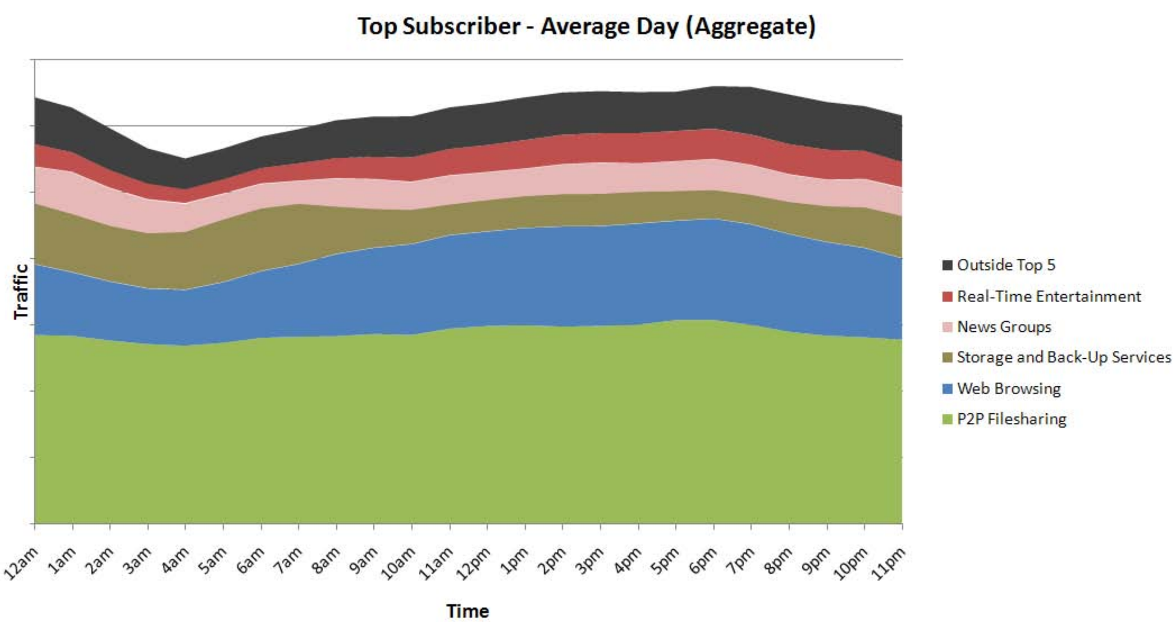

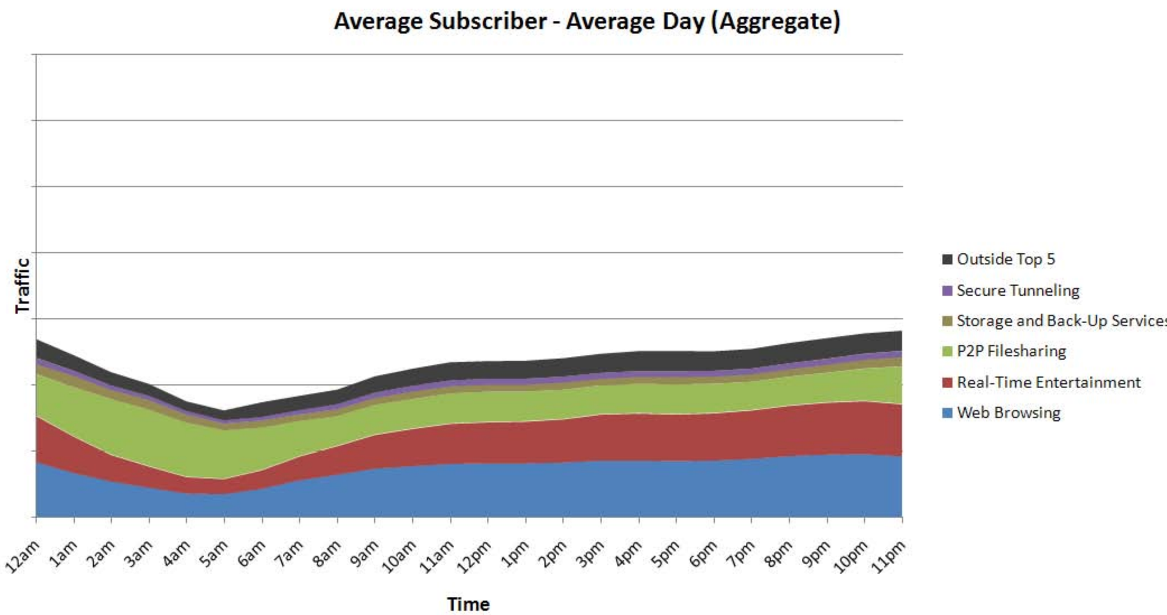

57. In contrast to a top subscriber (for whom P2P file-sharing accounts for almost half of all traffic), an average subscriber favours the on-demand nature of Web Browsing and Real-Time Entertainment. Also, while P2P file-sharing is still present in the average subscriber's profile, it accounts for less than a quarter of daily bytes.

Differentiated Treatment for Differentiated Demands

58. So, one important lesson to be learned from this data is that, left unmanaged, certain applications and subscribers win the inherent competition for shared network resources. To deliver the unique quality of experience that subscribers expect from each application for the maximum possible number of subscribers for the maximum amount of time, the network needs to differentiate between the heterogeneous needs of individual applications and subscribers. *Again: an unmanaged network is not a neutral network.* Not only should bright line rules prohibiting differentiation of network traffic be avoided, but differentiation should be encouraged in order to improve the experience of Internet users. The recent Canadian and European examples demonstrate growing global acceptance of this approach. As the Internet is a global, shared resource, global harmonization of regulation would be valuable goal for all Internet participants.

59. Sandvine is pleased to see that the Commission itself has acknowledged the benefit of differentiated treatment of network traffic on a per-subscriber basis. In his speech to the Brookings Institute, Commission Chairman Julius Genachowski noted, "During periods of network congestion, for example, it may be appropriate for providers to ensure that very heavy users do not crowd out everyone else."

60. Sandvine agrees with Chairman Genachowski's sentiments, and hopes that the Chairman and all Commissioners can apply the same logic to application traffic, so that no one category "crowds out" all others and so that all applications get the resources they need when they need them.

Differentiated Treatment of Network Traffic Achieves the Commission's Goals

61. If the Commission takes this step, it would move forward in achieving its goals for the NPRM. By helping each application get the resources it needs, application-specific differentiation of network traffic would help encourage investment and innovation in applications and related content, as it has done to date.

62. Application-specific network management practices also help further the Commission's goal in this NPRM to protect subscribers' rights. For example, in times of network congestion a subscriber-specific network management policy that "ensures that very heavy users do not crowd out everyone else" can be made less intrusive from the subscriber's perspective by layering on an application-specific component. By their nature, applications like VoIP, online video gaming and others do not demand much bandwidth and so do not contribute meaningfully to network congestion. However, because they are time-sensitive applications, their usefulness to the consumer is greatly impacted by any delays in their delivery. By combining a subscriber-specific and application-specific approach a network provider could create a narrowly-targeted policy that affects:
    - *only* disproportionate users;
    - *only* application classes that contribute disproportionately to bandwidth consumption; and
    - *only* application classes that are not time-sensitive.

63. In the future Sandvine expects to be able to offer solutions that let the *subscriber* select which applications receive higher priority in the network in times of congestion. By definition, any network management policy that is not only agreed to but *defined by* the Internet subscriber must be reasonable. A bright-line rule prohibiting differentiation between types of application-traffic could halt investment in this area, to the detriment of all Internet participants.

Innovation and Investment Continues to Occur at the Network Level

64. It is critical to remember that investment and innovation occurs not only by application and content providers, but by network providers as well. Thanks to past investments and innovation in network management solutions the Internet is already much more intelligent and capable of delivering a wider variety of services than at any previous time.

65. Network providers are just beginning to explore the use of network management practices to help them create service offerings that are more attractive to consumers in an increasingly competitive Internet access market. In the United States, high-speed Internet services are largely offered in the form of flat-rate, monthly 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 necessitate the ability to differentiate between the traffic of individual subscribers, and between applications.

66. For example, "light" Internet users may be interested in a service package that ties their fees to the bytes they consume on the network. But would these consumers want to pay for malicious traffic that affected their usage in a month, or visits to the service provider's web service portal to address service issues? A user- and application-specific policy would be required to manage the plan. By contrast, disproportionately heavy users likely don't want to pay "by the byte", but they may be interested in a service plan that provided a financial incentive to shift their activity to non-peak network hours. Such a plan would help all users by freeing up more capacity at peak times, when network congestion and application degradation is most likely to occur.

67. Other consumers may value their Internet connection by the quality of experience they receive for their favourite applications, like latency-sensitive Internet video gaming or VoIP. Network providers could offer a Premium Video Gaming or Premium VoIP service plan that delivers exactly the type of Internet experience these consumers want. Such plans would need to be supported by application-specific and user-specific policies.

68. By definition, any network management policy (whether user-specific, application-specific or both) that is deployed to support a term of a service plan that has been transparently disclosed and freely agreed to by both a network provider and an individual Internet subscriber must be deemed *reasonable*.

69. New service plans like these would represent significant innovations and require significant investments in network engineering (and marketing) by network providers. They would also offer consumers new choices and in so doing create new grounds for competition among network providers. Quite rightly, these are all goals of the NPRM.

The Internet is dynamic

70. The Internet is dynamic – this is another inescapable conclusion from Sandvine's 2009 Global Broadband Phenomena study. We are in the midst of a massive shift in subscriber behavior from a reliance on "download now, use later" content acquisition to an on-demand mentality where bytes are consumed as they arrive. The graph below depicts the percentage change between 2008 and 2009 in bandwidth share of various applications over a typical day and at peak hours only.

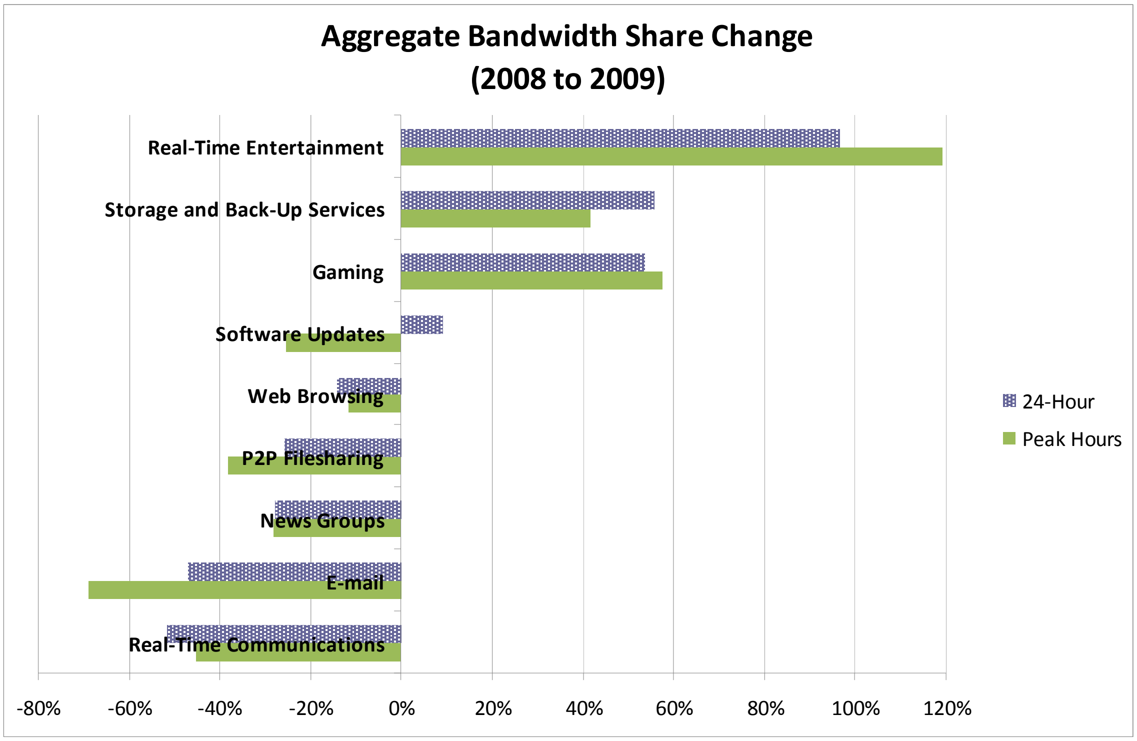

71. As described earlier, almost two-thirds of all Internet traffic in 2009 was enjoyed on arrival, including web browsing, real-time communications, gaming and real-time entertainment such as video and audio streaming and peercasting applications. The success of these applications throughout 2009 comes at the expense of traditional bulk data acquisition, most notably P2P file-sharing and news groups. During peak hours, when network congestion is greatest, the explosion in popularity of real-time entertainment and gaming applications has resulted in an even larger yearly increase than that witnessed in the 24-hour average.

72. In such a dynamic environment, it is even more important that network providers retain the flexibility to innovate with new network management practices that recognize the differentiated needs of applications – both those popular today and those as-yet-unknown applications with as-yet-unknown demands on network resources that will become popular tomorrow.

A Framework for Differentiation

73. Given the differentiated needs of applications and subscribers and the dynamic nature of Internet traffic, differentiated treatment of network traffic has been and will continue to be beneficial to the ongoing evolution of the Internet and consistent with the goals of the NPRM. As the Commission stated the problem in the NPRM, "The key issue we face is distinguishing socially beneficial discrimination from socially harmful discrimination in a workable manner." This statement itself recognizes that a bright-line rule prohibiting discrimination is inappropriate. There is good

discrimination and bad discrimination. What is needed is a framework to distinguish between the two.

74. As stated earlier in this submission: "Differentiation is good. Anti-competitive behaviour is bad." Sandvine submits that the first four rules proposed in the NPRM, plus the proposed "transparency" rule, all subject to "reasonable network management," provide a framework to allow for healthy differentiation and to prohibit anti-competitive practices.

75. Recognizing that the needs of the subscriber are paramount to the Internet, any Commission rules should be framed in terms of the user experience (as was the case with the four principles of the Commission's Open Internet Policy Statement) rather than as obligations of any individual Internet participant, such as network providers. Sandvine suggests rewording the first four proposed rules of the NPRM as follows:

    Subject to reasonable network management, consumers and providers of network access, applications and content should act in a manner that is consistent with the following rules:

    1. Consumers' are entitled to send or receive the lawful content of their choice over the Internet.

    2. Consumers' are entitled to run the lawful applications or use the lawful services of their choice.

    3. Consumers' are entitled to connect to and use on the network their choice of lawful devices that do not harm the network.

    4. Consumers' are entitled to competition among network providers, application providers, service providers, and content providers.

76. Under this framework, the definition of "reasonable network management" will play a pivotal role in defining healthy differentiation of network traffic. Sandvine's proposed definition is described in the related section below, under the heading "Reasonable Network Management."

Codifying a Principle of Transparency

77. In connection with the Commission's NPB Public Notice #24 (DA 09-2474) dated November 24, 2009, Sandvine made recommendations concerning how network and application providers could enhance consumer transparency of certain network performance metrics and application performance requirements, respectively. Refer to this filing for Sandvine's detailed comments on consumer transparency.

78. In the NPRM, the Commission is seeking comment on a more narrow topic, specifically "how broadband Internet access service providers should disclose

relevant network management practices to consumers as well as to content, application, and service providers and to government."

79. As with any other rules resulting from the NPRM, Sandvine submits that a transparency rule must be looked at strictly in relation to the needs of Internet subscribers. So, what do subscribers need to know? Subscribers only need information on any network management practices that affect in any way their entitlements (as described by the first four rules) to:
    - send or receive lawful content;
    - run lawful applications;
    - use lawful, harmless devices
    - competition among network providers.

80. If a network management practice does not affect their entitlements, no disclosure obligation should be triggered. If a network management practice does affect their entitlements, then network providers should disclose the potential impact on subscribers of such a practice in sufficient detail for consumers to be informed of the circumstances in which such entitlements may be affected and to what extent.

81. Sandvine therefore recommends that the fifth rule, the Transparency rule, be framed in a manner similar to the following.

    > Consumers' are entitled to disclosure of the pertinent details of any network management practices that affect consumers' entitlements as specified in this part in sufficient detail for consumers to be informed of the circumstances in which such entitlements may be affected and to what extent.

Reasonable Network Management

82. As with the framing of any rule resulting from the NPRM, Sandvine submits that the definition of reasonable network management must be framed in terms of the end user's Internet experience. Did the network management practice make the Internet experience better for most network users most of the time? Were users unreasonably limited in their access to content, applications, or devices of their choice as a result of the network management practice? Due to the heterogeneous and dynamic nature of the networks and traffic that comprise the Internet, such questions can only be analyzed and answers can only be provided on a case-by-case basis.

83. The Internet's constituent access networks, such as cable, DSL, fibre, wireless, and satellite, have different characteristics that create susceptibilities to congestion and performance issues at different network locations at different levels of usage. For example, in cable networks, upstream capacity has traditionally been very limited, in DSL networks certain downstream links can be more subject to congestion and in mobile networks upstream and downstream bandwidth is at a premium as it is fixed by the physical properties of the underlying radio spectrum.

84. Even within an access network category, no two service providers have networks that are identically architected. A particular technical approach that achieved reasonable results for users in one network may have a very different effect on users in another network – no bright lines can be drawn. Network management solutions are highly configurable. Both subscriber-specific and application-specific approaches can result in reasonable or unreasonable effects for users. The determination cannot be made ex-ante, only ex-post after seeing the actual effect on users.

85. Any definition of "reasonable network management" must therefore allow adequate flexibility to be adaptable to all network access technologies and all networks within each network access technology class. Sandvine recommends the following criteria for defining "reasonable network management".

- The network management policy serves a legitimate goal for that network, such as, but not limited to:
    - mitigating the effects on users of network congestion;
    - preserving the quality of experience for users' time-sensitive applications;
    - mitigating the impact to users of malicious or harmful traffic or content;
    - mitigating the unlawful transfer of content.
- The network management policy is tailored as narrowly as reasonably possible for that network:
    - to impact only users related to the policy goal, in network locations related to the policy goal, and only when needed to achieve the policy goal:
    - to achieve only the policy goal and not impact affected users in other ways;
- The network management policy is as minimally intrusive as reasonably possible for that network:
    - to differentiate between network traffic as little as reasonably possible to achieve its goal;
    - so that any disadvantage conferred by the policy to any end user's traffic is as minimal as reasonably possible
- The policy is auditable, so that network providers can demonstrate that the above conditions were met.

86. The CRTC in Canada arrived at a similar framework[12] in its October 2009 decision in its Review of the Internet traffic management practices of Internet service providers.

Applicability of Principles to Different Broadband Technology Platforms

- *Application of the Internet Principles to Wireless*
    - *Application of Nondiscrimination with Respect to Access to Content, Applications and Services, Subject to Reasonable Network Management the Internet Principles to Wireless*

[12] See http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm, paragraph 43

87. Regardless of the type of access network Sandvine recommends against a bright-line rule prohibiting differentiation between network traffic. This is especially the case in mobile networks, due to certain inherent characteristics.

88. Unlike fixed line networks, the bandwidth available in a wireless network is fixed and defined by its associated radio spectrum. With limited bandwidth, issues related to latency, jitter and packet loss also become exacerbated in the mobile environment. So, mobile networks are particularly susceptible to congestion and quality-of-service issues, and such limitations are already being noticed by users of some of the world's largest mobile networks[13] despite still-modest data usage. According to a July 2008 study by The Nielsen Company[14], only 15% of mobile subscribers in the U.S. actively used the mobile Internet. However, in a February 2009 study by the same company[15], 49% of U.S. non-users of the mobile Internet intend to use it in the next two years.

89. Mobile networks are also the newest entrants in the market for broadband access so user behaviour is rapidly evolving. Consequently, the nature of data traffic traversing the mobile network is more dynamic than for any other access network class. More time is required to understand how users will consume the Internet over mobile devices and what network management policies may be appropriate.

90. Sandvine has recommended that the Commission move forward with "user-centric" versions of the other four proposed rules, all subject to "reasonable network management", which should also be defined in terms of the user experience. If the Commission adopts this framework, Sandvine submits that all rules could be applied to mobile networks today, which would provide certainty for all stakeholders.

91. Since the standard for each rule would be based on what the *user* experiences on his network, Sandvine's framework would automatically allow for *different network management practices based on different network characteristics*.

92. For example, a few bulk file-sharing or file-transfer sessions is unlikely to have a crippling effect on users' enjoyment of other applications in today's fixed line networks. The same may not be true for a given mobile network. A few similar sessions on a cell site could seriously impair the web surfing, voice call and gaming experience of all users' connected to that site, for example. Consequently, to protect the user experience for these popular applications, it would be reasonable to create a policy that began to manage bulk application traffic at a lower threshold (and/or manage it in a different way) than for a fixed line network. In fact, these applications, and others like Slingbox (which "slings" bandwidth-intensive television signals to

---

[13] ArsTechnica. *AT&T CTO downplays role of iPhone in network's issues*. See http://arstechnica.com/apple/news/2009/10/att-cto-downplays-role-of-iphone-in-networks-issues.ars?utm_source=rss&utm_medium=rss&utm_campaign=rss

[14] The Nielsen Company. *Critical Mass. The Worldwide State of the Mobile Web*. See http://www.nielsenmobile.com/documents/CriticalMass.pdf

[15] See http://www.tellabs.com/news/2009/index.cfm/nr/53.cfm

Internet devices, such as a Smartphone), may have such a detrimental effect on network performance for all applications that blocking them could be deemed a reasonable practice in a given situation at a given time. A case-by-case analysis would have to be performed to know.

93. Similarly, if managing mobile data traffic on a subscriber-specific basis, it might be necessary to start managing "disproportionate users" consumption at a threshold level that would be much lower (or managing it in a different way) than for fixed line networks.

94. Sandvine believes in a common set of rules for all users in the Internet commons. Sandvine's proposed rule framework, including its definition of "reasonable network management", allows for this to happen today by providing the Commission a single method for analysing network management practices across all access networks.

Appendix 1 – Assumptions for Application Requirements

95. P2P, Usenet, FTP: Bandwidth is the most important network characteristic as it affects the time required for these applications to transfer the data. A typical movie is about 700 MB. If a typical user expects, at a minimum, to download a movie overnight (i.e., 8 hours), the minimum bandwidth required would be 195 Kbps.

96. Web Surfing: A typical "Web 2.0" website requires approximately 10 to 20 connections to download the approximately 0.5MB to 2MB of data needed to display the page. Studies have shown that to maintain a good user experience this must be done within 2-4 seconds[16].

97. To reach a website, its name must first be translated into its numeric IP address via the DNS. This happens for each server that the webpage references. Many webpages have images, videos and advertisements on different servers and thus the Internet browser must resolve each DNS name. Each time DNS is used, 2x the latency (for the round trip) is added to the total time to load the page. Also, to compensate for jitter in most PC environments, the PC buffers the data to the extent of the jitter in the network so that the total latency time is actually (latency + jitter).

98. For each connection or file that needs to be downloaded, the (latency + jitter) is added multiple times as the browser initiates a TCP connection to the server to retrieve the file. This multiple is usually three, in accordance with TCP's three-way handshake for initiating connections.

99. Once the connection is established, the time to download the file is a function of the bandwidth available. However, given that websites often have many small files (images, text), TCP is not always able to achieve the full throughput rate due to its "slow-start[17]" algorithm.

100. The above argument does not take into account many of the complex algorithms or tools in place such as parallel connections and HTTP pipelining, but does show that bandwidth is not the only determining factor for measuring HTTP quality of experience. In fact, latency and jitter will likely be the gating factors on the user's quality of experience with Web 2.0 websites.

101. For a 1.5MB webpage with 20 connections to load with a satisfactory user experience, available bandwidth must be at least 1 Mbps and latency+jitter must not exceed 166ms.

102. Email: A normal text email is between a few kilobytes and a few hundred kilobytes. Email is not instantaneous, however, there is a perception that it is near real-time. To send or receive an email with a large attachment in under a minute, the bandwidth required is approximately 60Kbps.

---

16 See http://www.akamai.com/html/about/press/releases/2009/press_091409.html

17 IETF RFC 2001. See http://tools.ietf.org/html/rfc2001#ref-2

103. VoIP: The most basic audio codecs require bandwidth of approximately 16 Kbps (allowing for overhead of the Internet). VoIP is a real-time application that is very sensitive to latency and jitter. ITU-T G.114 suggests that the maximum one-way latency+ jitter be 150 ms18 (or round-trip 300 ms), above which it becomes noticeable to the end user. Most VoIP protocols use stateless connections (UDP) and have no built in retransmit. Loss must not be over 0.5% for calls to be audible.

104. Video Conferencing: Similar to VoIP, the application is bi-directional and is highly susceptible to latency, jitter and loss, however the bandwidth requirements are higher due to the addition of the video.

105. Video & Audio Streaming. These applications are primarily uni-directional. The average normal-definition video on YouTube requires approximately 300Kbps. High-definition videos (depending on the quality, i.e., different encodings) require bandwidth between 1-3Mbps. Because most streaming done on websites like YouTube use HTTP, latency, jitter and loss are not a major concern. Traditional streaming video (RTSP, RTP, etc) are done over UDP and are affected more by loss. Streaming of compressed CD quality audio requires approximately 160Kbps of bandwidth.

18 International Telecommunication Union. ITU-T Recommendation. G.114. See http://www1.cs.columbia.edu/~andreaf/new/documents/other/T-REC-G.114-200305.pdf