October 29, 2009

VIA ELECTRONIC FILING

Marlene H. Dortch
Secretary
Federal Communications Commission
445 12th Street, S.W.
Washington, D.C. 20554

Re: In the Matter of Broadband Industry Practices, WC Docket No. 07-52, and GN Docket 09-191.

Dear Ms. Dortch:

Pursuant to Section 1.1206 of the Commission's rules, I, George Ou, hereby give notice that on October 2009, I met with Christi Shewman with regard to the above-referenced matter. This meeting was conducted in offices of with Commissioner Meredith Attwell Baker. The parties discussed the attached materials.

The purpose of the meeting was to explain the importance of "reasonable discrimination" based on an article published at Digital Society[i]. The contents of the article have been attached to the end of this document.

The current 5th principle of the NPRM prohibits any sort of discrimination in favor of or against any particular applications over other applications. And while the "reasonable network management" exception mentions the possibility of prioritizing some applications such as VoIP, it raised concerns about the fairness of such a scheme and it sought comments from the public.

What I explained that was that the FCC already allows Comcast to "discriminate" in favor of low bandwidth subscribers over high bandwidth subscribers because that is the fair thing to do. But the same concept should also applies to applications and low bandwidth applications such as VoIP or online gaming should always be prioritized over high bandwidth applications like web browsing. Since Web browsing is a low duration application i.e., it bursts a lot of bandwidth but quickly drops to zero when the web page isn't loading, web browsing should always be prioritized over high bandwidth and high duration applications like file transfer e.g., BitTorrent or other Peer-to-Peer (P2P) applications.

The second purpose of the meeting was to discuss the importance of having ample competition in the content distribution space. Unfortunately in the current draft of the NPRM in section 106, this new rule would essentially cut down competition in the content distribution space by specifically singling out the Broadband industry and prohibiting them from offering "premium" or "prioritized" services. That would presumably prohibit edge caching services or priority bandwidth or priority queuing though the specific language in the NPRM did not specify.

While section 106 may be well intentioned and based on what seemed to be an ultimatum from former SBC and AT&T CEO Ed Whitacre that application and content providers on the Internet weren't going to use his pipes for free, the approach in section 106 is too heavy handed. Whitacre's comments were perceived to imply something sinister, that AT&T would presumably block or degrade customers unless application and content providers paid AT&T for premium or prioritized service. The theory was that AT&T and other broadband providers would use their market power to double dip by charging the broadband subscriber and then recharging the content provider for same bandwidth that the broadband subscriber already paid for.

Yet over the course of the last 3 years, this threat never materialized and the opposite actually happened. It turned out that network providers didn't have this type of market power to intimidate content or application providers but ESPN could exploited their market power by blocking entire (Internet Service Provider) ISP networks from accessing ESPN360 until the ISP paid a per-subscriber fee to ESPN[ii].

What did happen was that content and application providers like Amazon actually went to network operators like Sprint and offered to pay for mobile Internet service on the Amazon Kindle eBook reader on behalf of the customer. That means the customer wouldn't need to pay anything for mobile Internet service because Amazon was able to negotiate a better group rate for all Kindle users. Yet the new FCC NPRM in its current form would prohibit Broadband providers from charging Application or Content providers which mean America would lose these innovative devices and business models like the Kindle.

A more precise rule with less collateral damage would be to simply prohibit double dipping on the part of the broadband provider. In other words if the customer already paid for a certain level of bandwidth and certain level of priority, ISPs should not be allowed to charge content providers again for the same level of service. But this should not preclude an ISP from charging application and content providers for bandwidth or priority that the broadband subscriber never paid for. Moreover, it should not preclude the ISP form offering edge caching services and competing in the content delivery market.

Some would argue that ISPs are too powerful of a "gatekeeper" and that they should not be permitted to discriminate, but section 106 goes so far as to declare the mere offering of content delivery services as an illegal form of discrimination. But the idea that the ISPs are gatekeepers is outdated because more and more high quality content is bypassing what we consider the traditional Internet and moving onto private network backbones and private server caching infrastructure.

According to a new Arbor Networks report, Google's traffic (largely YouTube) is approaching 10% of total Internet's traffic. Google is its own network backbone provider that has the size and scale to directly peer with broadband providers[iii]. That means Google is now one of the world's largest network operators that has its own private Internet backbone and massive server farms that are deployed across strategic places across the globe. Google has now emerged as one of the new gatekeepers and it is quickly obsolescing the old traditional Internet gatekeepers.

Yet because of some regulatory quirk in the new FCC NPRM which singles out broadband providers and prohibits them from competing in the content distribution space, Google the networking and server farm giant could effectively be granted a monopoly by the FCC if the NPRM section 106 is allowed to stand in its current form. The effect of NPRM section 106 would be less competition in the content distribution space and higher costs for content producers, which means the additional cost would ultimately be passed on to consumers.

FCC 5th principle must allow for reasonable discrimination

BY GEORGE OU 22 SEPTEMBER 2009

Before I start this discussion on the newly announced 5th FCC principle of the Internet, I want to explain my use of the word "discrimination" in this article. The word discrimination over the years has taken a different meaning from the classical definition of the ability to differentiate. It has become synonymous with the immoral and inexcusable practice of prejudice against a certain group of people. Most recently, "discrimination" is now commonly used to describe unethical prejudice against applications, users, and businesses of the Internet. The problem with this evil-only definition of the word "discrimination" is that it forecloses the possibility of reasonable discrimination when the original definition had the ability convey both good and bad differentiation. So for the purpose of this article, I'm going to use the classical generic definition of discrimination which does not infer good or bad and I will instead specify what reasonable discrimination is and what is not.

Introduction

As most of us have heard by now, the FCC announced yesterday their intention to create a new "5th principle" of nondiscrimination on the Internet. While there are some serious questions as to whether this is a wise mandate for wireless Internet services, such a principle applied to wired Internet services could be a good thing if it allows for reasonable discrimination. Reasonable discrimination may fall under the following two categories.

Reasonable network management

Reasonable business practice

Reasonable network management

The goal of network management is to ensure proper and efficient sharing of network capacity. All networks have to slow down users and applications when there is simultaneous activity, but the question is how much each user and each application is slowed down relative to other users and applications. It turns out that an unmanaged network is the least fair of all because the least aggressive users and applications are stomped into the ground. I published a detailed report last year on reasonable network management which offers additional insights into these issues. I also have a 7 minute animated presentation here that explains the need for a more intelligent prioritized network.

Some aggressive users and their applications can consume 10 to 40 times more bandwidth than other users in a congested network. When usage duration is factored in, those same aggressive users can account for 100 to 400 times the network usage compared to the typical network user. It is reasonable in this case to deprioritize these aggressive users behind users who are asking for only a tiny fraction of the total capacity. This is reasonable discrimination intended to counter a much more harmful unfairness of TCP that allows aggressive users and applications to suppress other users and applications. The FCC has already reviewed this type of prioritization scheme from Comcast's new "fair share" network management system and they had no issues with it.

Some applications like Voice over IP (VoIP) or web browsing have very low average bandwidth consumption. The VoIP application uses a slow and steady stream of bandwidth while the web browser might an occasional sharp burst in bandwidth when they load new web pages but have very low average bandwidth requirements because the browser spends most of its time idle as the user is reading the content. Low bandwidth applications like VoIP might typically only need 1% of the network's capacity but they cannot tolerate any kind of delay because the real-time nature of voice communications means that delayed transmissions have to be discarded.

Peer-to-peer (P2P) applications will aggressively consume 100% of any available bandwidth by opening up dozens of communication channels called "flows" and they will consume 100% of the

duration by remaining constantly on. Even when a P2P application is configured to only use 10% of the network's capacity, it can still cause severe problems for VoIP because it has a tendency to burst traffic and cause a micro-congestion storm where 100% of the network's capacity is filled for a fraction of a second or more. This micro-congestion storm is called "jitter" and it has the ability to completely block a VoIP conversation for a fraction of a second or more which can cause one or several words to go unheard. In this case, it is reasonable to always prioritize the low bandwidth VoIP application over the high bandwidth P2P application.

Web browsers are aggressive with bandwidth by opening up to 4 simultaneous flows, but they only ask for bandwidth for a few seconds and then they go idle for several seconds or minutes while the user reads the content of the webpage or does something else. The web browser can cause small amounts of jitter and for this reason, it makes sense to deprioritize web browsing traffic behind VoIP applications. But because it is a low duration application and because it is far less aggressive than P2P applications, it makes sense to prioritize web browsing traffic ahead of P2P traffic. Furthermore, the beauty of this scheme is that it does not increase the overall download time for P2P traffic. That's because higher bandwidth priority for the short duration web browsing traffic allows it to complete sooner which allows the P2P application to resume full speed sooner. The result is no change in the performance of the P2P application but a vastly improved web browsing experience. See figure 1 below.

Figure 1

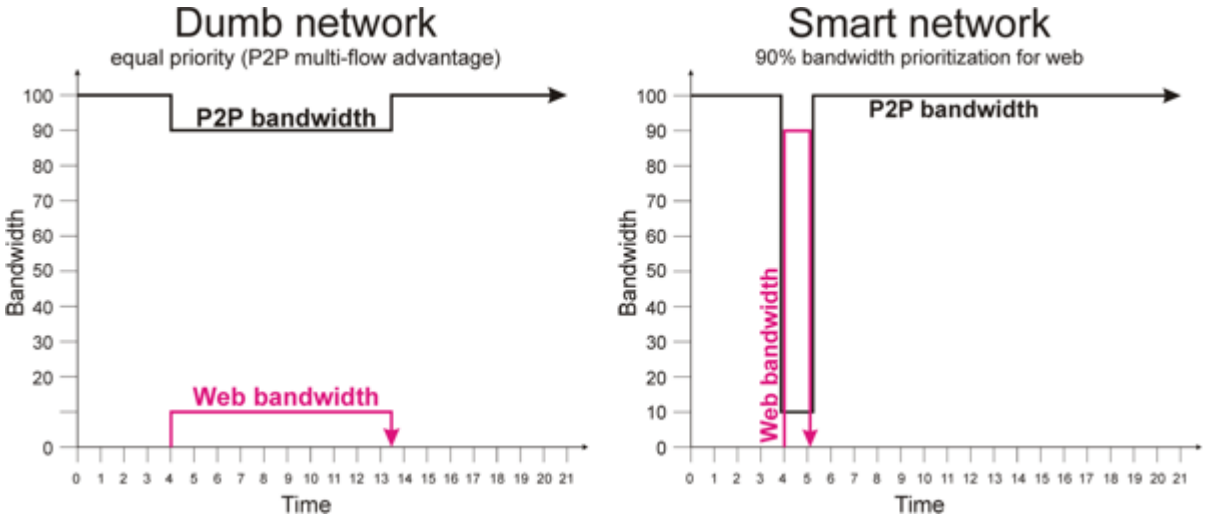

Furthermore, because the P2P application is no longer "toxic" to other applications like VoIP or online gaming that the user may want to simultaneously run, the users no longer needs to stop or severely throttle their own P2P usage to accommodate other applications. So by deprioritizing P2P with the smart network, we actually improve P2P performance.

Universal rule for reasonable application priority

Based on these facts, we can create a sort of universal rule for a reasonable network management scheme. Any prioritization scheme that prioritizes low bandwidth application ahead of a high bandwidth application, and any low duration application ahead of high duration application is reasonable and actually beneficial to the applications given low priority. Figure 2 shows the reasonable order of packet priority for various types of applications.

Figure 2

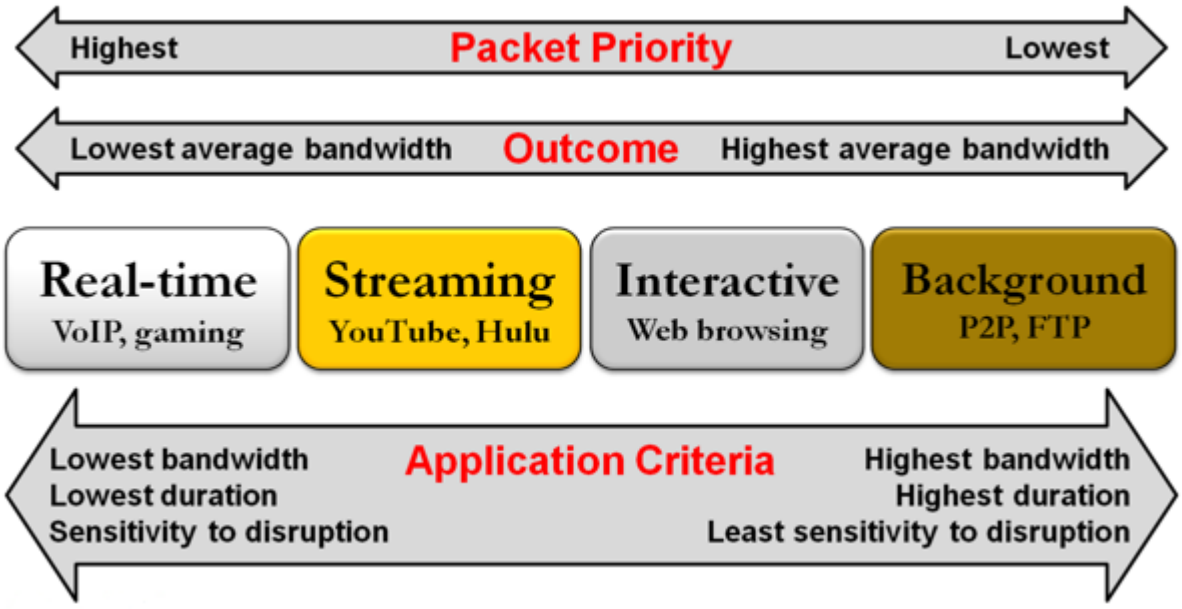

So using this guideline, the FCC has at its disposal a “litmus test” for determining what is reasonable discrimination and what is not. If an ISP classifies a low bandwidth VoIP application as a background low priority application, that would obviously be unreasonable discrimination.

The VoIP Company Skype has publically raised concerns about any prioritization scheme since they like to call themselves a “P2P” application and fear that they would be misclassified as a low priority application. However, this has never been the case since Skype is a low-bandwidth application which would automatically classify it as a high priority application regardless of the “P2P” classification given to it by Skype. It’s also important to note that Skype only uses P2P as a protocol of last resort when all other methods of working around network firewalls and routers fail, so the P2P label is questionable to begin with since it rarely resorts to P2P mode.

It is also important to note that the web browser can be used as a high duration large file transfer mechanism so in that particular instance, the web traffic would actually be classified as a background low priority application. The key factor in determining priority is the behavior of the data flow so a single application could have multiple classifications depending on how it is used.

Should ISPs or their users determine priority levels

Some have argued that only the end user should determine the priority of their own packet priority, but banning ISP involvement would be unworkable in the real world. The first problem is that network prioritization works best on the transmit end. That means downstream (download) traffic is always best handled by the ISP. The end user can do some limited amount of management on the receive end, but the result is poor throughput for the high bandwidth application and minimal jitter mitigation which results in poor VoIP or online gaming performance. Managing the network on the transmit end allows the high bandwidth application to run at maximum speeds while completely eliminating the jitter problem for optimum VoIP and gaming performance.

The other challenge for a user-only scheme is what happens when the user labels 100% of their packets from every application as high priority? Some popular P2P applications even allow users to set their own packet priority and this would break the entire network management scheme. Business class ISPs have allowed their business subscribers to choose their own priority levels for their own applications, but there are contractual limits on the various priority levels based on what the customer is willing to pay are enforced. If the customer exhausts their priority budget, then all of their subsequent traffic is treated as background priority regardless of the priority label.

If an ISP enforces priority budgets, it would technically be possible to allow the end user or application to set their own priority levels and override the ISP’s priority scheme. The problem with this is that the vast majority of users lack the knowhow or the desire to become network engineers since computers

and home networking is already complicated enough. Even in the case of business subscribers who routinely set their own priority levels, this task is handled by the network engineers employed by the business. Since residential broadband customers usually don't employ their own network engineers, the only way that the vast majority of users can benefit from intelligently managed networks is if the ISP's network engineers manages the priority for them. While some would raise the possibility of ISP abuse, the fear is unfounded so long as the ISP follows the universal guidelines for packet prioritization in figure 2.

Reasonable business practices

While figure 2 is a great guideline for prioritization schemes, there are exceptions that may pop up under reasonable and existing business practices. To encourage private investments in a free society, businesses must have the right to use private property and private capacity to earn a return on their investments. If an ISP prioritizes subscription television services like IPTV by reserving a fraction of the total broadband capacity despite the fact that it is a high bandwidth and high duration application, this is a reasonable business practice and it is used all over the world. Furthermore, letting telephone companies offer IPTV to compete with cable TV encourages investment in higher capacity Telco broadband which in turn spurs cable TV companies to invest more in faster cable broadband.

Regardless of what some Internet companies who don't have to spend billions of their own dollars in building the broadband networks may think, it is reasonable for a Telco to dynamically set aside a portion of the physical network infrastructure they invested heavily in to offer services like telephone or television service in addition to broadband services. This is no different from a private company building private communications capacity and we respect these property rights in a free society. This is no different than a cable TV company statically setting aside 95% of its physical coax cable infrastructure to television services and the remaining for Internet capacity. The difference here is that dynamic bandwidth allocation used in IPTV allows the consumer to reclaim their television bandwidth for Internet access when they choose to turn the television off. Without IPTV prioritization, Internet activity can easily disrupt television services which consumers simply won't tolerate. Not prioritizing IPTV would force consumers to choose between using the Internet and watching IPTV but not both at the same time. Having the IPTV prioritization in place typically only consumes a small percentage of the total broadband capacity and it allows the peaceful coexistence between IPTV and broadband usage.

Some might argue that broadband providers should no longer be allowed to offer traditional subscription television services, but this is economically untenable so long as we expect private investment to build the next generation broadband network. Even publically run municipal broadband companies have had to rely on television subscriptions services that reserve network capacity to stay afloat financially so why would anyone expect a private company to be any different? So long as television services remain a critical revenue stream that supports the expansion of broadband capacity, it must be permitted to exist as a reasonable business practice.

---

[i] George Ou, Digital Society, September 22, 2009, "FCC 5th Principle must allow for reasonable discrimination", http://www.digitalsociety.org/2009/09/fcc-5th-principle-must-allow-for-reasonable-discrimination/

[ii] George Ou, ZDNet, June 21, 2006, "Is ESPN committing reverse Net Neutrality?", http://blogs.zdnet.com/Ou/?p=251

[iii] Ryan Singel, Wired, October 16, 2009, "YouTube's Bandwidth Bill Is Zero. Welcome to the New Net", http://www.wired.com/epicenter/2009/10/youtube-bandwidth/