

**Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, D.C. 20554**

In the Matter of

Broadband Practices

)
)
)
)
)

WC Docket No. 07-52

SECOND REPLY TO “TESTIMONY OF GEORGE OU”

Regarding “Comments on inaccurate testimony at the FCC broadband Industry Practices
Hearing;
Stanford University;
April 28, 2008”

Robert M. “Robb” Topolski
6815 NE Vinings Way Apt 922
Hillsboro, OR 97124
(503) 342-2468

The following replies were made to George Ou’s blog post at <http://www.formortals.com/Home/tabid/36/EntryID/11/Default.aspx> which was filed in ECFS as http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6520001501. My reply was filed as http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6520004355 but the dialog, unfortunately or not, did not stop there.

This filing contains and updates the information filed in http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6520004355. Modifications have been made only as to document form. I have made no modifications to the text except where marked *-rmt-*.

As the attempted dialog is going nowhere, except to provide a platform for personal and/or professional attack, I have ended my participation in it for now and deny that anything has been “debunked” in either direction. My lack of participation in this dialog going forward is not an indication of agreement, disagreement, or lack of a different view.

Thank you for reading and I apologize in advance for any contribution that I may have made to the discouraging behavior the reader is about to witness...

Comments on inaccurate testimony at the FCC Stanford hearing

Apr 27

Written by: George Ou
4/27/2008 4:56 PM

The FCC hearing at Stanford University on April 17th 2008 was filled with inaccurate testimony from various witnesses. Since that testimony seems to be carrying significant weight both on Capitol Hill and in the media, I feel compelled to set the record straight. I have filed a [copy of this letter](#) on FCC docket 07-52.

Problems with Jon Peha's testimony

Jon Peha testified that BitTorrent was like a telephone and implied that if a TCP reset is used by Comcast to stop a TCP stream, then that constituted a blockage of BitTorrent. Furthermore, Professor Peha implied through his telephone analogy that if BitTorrent is blocked, then the user must manually redial to reestablish the connection. These assertions are highly inaccurate and here's why.

The first problem is that Jon Peha did not understand the multi-stream aspect of BitTorrent or P2P. Peha seemed very surprised immediately before our panel at Stanford when I told him that a P2P download typically used 10 to 30 TCP streams at the same time. His surprised reply to me was "all active?" and I replied yes. The reality is that if a certain percentage of BitTorrent TCP streams are reset and temporarily blocked by an ISP, say 15% for example¹, then the "Torrent" (the file that's being exchanged amongst multiple peers over the BitTorrent protocol) is essentially slowed down by an average of 15%. In other words, the "Torrent" would suffer a 15% partial blockage which is accurately described as a "delay" since the file transfer didn't actually stop. This would be like filling up a bath tub with 20 faucets and you closed 3 of those faucets. The rate of water flowing in to the tub would slow but not stop.

The second problem with Jon Peha's testimony is his implication that the user must take some sort of action to resume the BitTorrent connection or else the connection wouldn't resume. Peha's assertion can easily be proven false by a simple experiment with BitTorrent. One can easily confirm that BitTorrent will always resume a lost connection within a matter of seconds without any user intervention just by physically disconnecting a network cable on the test machine and reconnecting it. Not only does BitTorrent automatically resume, it picks up where it left off and does not need to start all over again. So even if all TCP streams in a Torrent were simultaneously blocked for a short period of time, it will quickly resume by itself and eventually finish the file transfer. Therefore this is by definition a "delay" and not a "blockage".

This is not to say that Comcast's existing form of network management is without problem because it is clear that the system has flaws and unintended consequences like the accidental blockage of IBM Lotus Notes. The use of TCP resets also have a more drastic effect on rare Torrents, which are BitTorrent files that are not popular and have few seeders or other peers who have parts of the file to download from. These rare Torrents aren't healthy to begin with and a TCP reset can in some cases trigger a complete temporary blockage. The rare Torrent will still get there eventually but it will suffer significantly more than a normal Torrent that is resilient to partial blockage.

It should be noted that BitTorrent in general is not an appropriate file transfer protocol for rare Torrents. BitTorrent tracker sites tend to rank and sort a Torrent based on the Torrent's "health" which is based on the number of available seeders and pre-seed peers. Users generally tend to avoid the "unhealthy" Torrents on the bottom of the list. Since Comcast offers a vastly superior alternative where they provide 1 gigabyte of web storage space, Comcast customers can use that service to distribute files 10 to 20 times faster than any

single Comcast BitTorrent seeder could ever provide. To further illustrate this point, Richard Bennett [posted a copy of the King James Bible](#) on his Comcast-provided web space. By contrast, I couldn't find a non-copyrighted version of the King James Bible or any non-copyrighted Barbershop music on any BitTorrent tracker site which illustrates how unlikely it is that you'll find rare or legal content on BitTorrent.

Problems with Robert Topolski's testimony

Robert Topolski also had problems in his testimony. Topolski, a software tester who does not work in the networking field, insists that the TCP reset mechanism isn't common in network devices and declared that I was wrong in my testimony. In my experience as a Network Engineer who designed and built networks for Fortune 100 companies, it is my experience that the TCP reset mechanism is common in routers and firewalls. For many years, Internet service providers, including LARIAT (owned and operated by Brett Glass, who has filed comments in this docket) have used TCP RST packets to **protect the privacy** of dialup Internet users. When a dialup user's call is over, RST packets are sent to terminate any remaining connections. This prevents a subsequent caller from receiving information -- some of which might be confidential -- that was intended for the caller who disconnected. Thus, the transmission of RST packets by a device other than the one(s) which established a connection -- for the purpose of informing the endpoints that the connection has been terminated -- is not only commonplace but salutary. Network architect Richard Bennett who works for a Router maker explained to me that TCP resets are the standard mechanism used by consumer routers to deal with NAT table overflow, which itself is typically caused by excessive P2P connections. Are we to believe a single software tester or three networking experts?

The second key problem with Topolski's testimony is that to my knowledge, he has never provided any forensic data from Comcast in the form of packet captures that can be independently analyzed. Even if Topolski did produce packet captures and we assumed that those packet captures are authentic, one man's packet captures wouldn't be a large enough sample to draw any conclusions by any legal or scientific standards. The Vuze data may constitute a large enough sample but the data isn't very granular because it doesn't tell us what percentage of reset TCP sessions are due to an ISP versus other possible sources¹. The Vuze data, which even Vuze admits isn't conclusive, is highly suspect because it even shows as much as 14% resets on 10-minute TCP sessions from ISPs who do NOT use TCP reset packets to manage their network.

Furthermore, even if a TCP reset was used by Comcast at 1:45AM, we cannot assume that there was no spike in congestion at 1:45AM. As I have indicated in the past, just 26 fulltime BitTorrent seeders in a neighborhood of 200 to 400 users can consume all of the available upstream capacity in a DOCSIS 1.1 Cable Broadband network. That means less than 10% of the population seeding all day and night can cause congestion at any time of the day. Based on what little evidence presented by Robb Topolski, no conclusions can be drawn regarding the question of whether Comcast uses TCP resets for purposes other than congestion management.

1. The reason I use 15% as the example is because the Vuze data gathered via thousands of user's computers indicated that a Comcast broadband network typically suffered 14% to 23% for all TCP streams during 10-minute sampling periods. That 23% figure is not restricted to just BitTorrent or P2P traffic and even those TCP resets pertaining to BitTorrent aren't necessarily from Comcast. It's quite possible that the reset actually came from the client on the other end or it may have come from a customer-premise router on either end-point trying to conserve on NAT (Network Address Translation) resources. It is undeniable that a certain percentage the TCP reset have nothing to do with Comcast's network management practices. We really can't know what percentage of those TCP resets were due to Comcast and figuring out the exact percentage isn't trivial because there are so many factors to consider.

Note: Any assertions on behalf of Brett Glass and Richard Bennett that I have made in this document have been approved by Brett Glass and Richard Bennett.

20 comments follow ... -rmt-

Re: Comments on (mostly)-accurate testimony at the FCC Stanford hearing

Dear George,

I noticed that you consulted with Brett Glass and Richard Bennett before posting your article. You certainly have my permission to consult with me, too. Why should we be on completely opposing sides of this debate? You know how to find my phone number.

FIRST, AN APOLOGY: At the Stanford hearing, you reported that some smart gateway/router devices know when a host behind it has gone offline, and will issue RSTs in lieu of forwarding to the known-offline host. At the hearing, I misunderstood what you were saying and I countered that you were describing behavior covered in the standards. While I don't know of any particular smart devices that do this, if they do (and it makes sense to me that they do so), then I was incorrect to counter you without acknowledging that gateway devices might do this exactly as you described. For that, I **WHOLEHEARTEDLY APOLOGIZE**. Had I understood your statement (which is clear enough in the video --- the fault of misunderstanding is all mine), the behavior you described is probably not in the standards (AFAIK). However, the device sending the RST is acting as the end point and was instructed to do so by some administrator. It was not a secret. It was not DPI changing the behavior of the Internet. As far as the Internet peers were concerned, the SYN-RST exchange was completely expected. The entire discussion ended appropriately when Professor Peha, countering both of us, made the entire discussion moot when he pointed out that these are not examples of using RST for congestion control. You weren't wrong in what you said. I **WAS WRONG** to dismiss it as errant and I apologize for doing so.

SECOND: "P2P download typically used 10 to 30 TCP streams at the same time." This is essentially correct, but moot. Also the statement is true only for **DOWNLOADING WITH BITTORRENT** specifically. It is not true for eMule or Gnutella. It is not true while BitTorrent is only uploading. While partly true, the statement is completely unimportant. The important direction is the upload direction -- because the transmitting host handles congestion control on the Internet, because the last-mile congestion is the congestion we are discussing (this is the uploader's first mile), and because P2P uploads is what Comcast is tearing down. With BitTorrent, only 3-4 streams are allowed to upload simultaneously. This is a major flaw in **YOUR** testimony to the FCC and other bodies (I'm referring to your article and graphic showing a single connection comparison to 11 connections). Limiting the uploading connections to 3-4 is a specific feature of the BitTorrent protocol designed to prevent congestion by quickly responding to changing network conditions. See http://wiki.theory.org/BitTorrentSpecification#Choking_and_Optimistic_Unchoking to understand how this feature is designed to prevent congestion.

THIRD: "The second problem with Jon Peha's testimony is his implication that the user must take some sort of action to resume the BitTorrent connection or else the connection wouldn't resume." Jon Peha is entirely correct, as long as you understand "user" as being the application and not the person. The response to a RST action is completely dependent upon the application programming. An application developer has to choose whether to retry an RST-aborted connection, and when. The scenario is not covered in the BitTorrent protocol, so the developer has free reign to drop the peer, retry it immediately, or drop the peer from the peer list. The behavior is particularly damaging to users of current versions of eMule and Shareaza (which supports Gnutella and Gnutella 2) as a client who resets is penalized by those applications. With eMule and Shareaza peers, you lose your place in line and have to wait at the back of the queue. With Shareaza, the affected peer is also marked as "suspicious" on the network until it is able to complete an upload transfer to someone.

FOURTH: 'So even if all TCP streams in a Torrent were simultaneously blocked for a short period of time, it will quickly resume by itself and eventually finish the file transfer. Therefore this is by definition a "delay" and not a "blockage".' **YOUR STATEMENT IS NOT TRUE** when the affected peer is trying to upload unique content, such as their own musical tracks or movies. When the Associated Press and EFF created new BitTorrent uploads and tried to transfer them over Comcast links, the transfers were prevented. A disconnected TCP session cannot carry data. I have some responsibility for the word, "delay" as I used the word when I first posted about this on

DSLReports.com. (I also, unfortunately, used the word "manage," too.) As it turns out, both of those words proved to be too generous.

FIFTH: In your article today, the very next paragraph recognizes the problem. You said, "The use of TCP resets also have a more drastic effect on rare Torrents, which are BitTorrent files that are not popular and have few seeders or other peers who have parts of the file to download from. These rare Torrents aren't healthy to begin with and a TCP reset can in some cases trigger a complete temporary blockage." I think we're close to agreement, here. However, being the lone seeder of my own unique content is EXPECTED. The swarm cannot become healthy if that lone seeder is prevented from uploading. This was the case for me with Gnutella. Until February 20th or so, I always have enjoyed some ability to upload with BitTorrent, despite considerable interference. As of my last tests, I cannot seed at all. Because new torrents are rare torrents until sufficiently seeded, your sentence "It should be noted that BitTorrent in general is not an appropriate file transfer protocol for rare Torrents," is obviously false. As an experiment, I uploaded the Public Domain Gutenberg version of DaVinci's Notebook to The Pirate Bay last July. Despite not being Moviez, Tunez, or Warez, the swarm started with only me, and grew to hundreds large. A single uploader is quite important, and P2P introduced and maintained the availability of this content for months. I've long since deleted my copy, but the file remains available on TPB today! See http://thepiratebay.org/tor/3735336/The_Notebooks_of_Leonardo_Da_Vinci_Complete_by_Leonardo_da_Vin and download it, should you want.

SIXTH: "Topolski, a software tester who does not work in the networking field," George, at age 20, I was on the air with both RTTY, Fax, Slow-Scan TV, and etc.. These are point-to-point protocols, to be sure. By 24, I was heavily into store-and-forward systems. Law Enforcement was my profession then, but digital communications was my hobby. It would become my career in my 30s. Except for about 3 years, all of my tech jobs have involved networking. Desqview/X, QMosaic, Quarterdeck WebAuthor and WebServer, Intel VideoPhone (H.324), CNN at Work, Intel ProShare Presenter, the internal conversion to IE6, a public-facing and private-facing Y2K database application, an internal and external Learning Management System, and Intel Servers were a few of my projects during my career. Protocols come natural to me because I started studying them while data rates were slow enough that it was still possible to decode common ASCII phrases by ear! Companies don't hire people that know things, they hire people that DO things. That's why my job titles have been Quality Assurance and Testing related, but one of my top strengths, obviously, is networking.

SEVENTH: "[Topolski] insists that the TCP reset mechanism isn't common in network devices and declared that I was wrong in my testimony." I stand by my statement. They are not COMMON in end-user network devices, although some exist. I've used routing and NAT devices from D-Link, Microsoft, GlobalSun, Unicom, Linksys, Westell, Netgear and ZyXEL. None send resets (even if they can and should) -- they drop the packet instead. I said that they were not common on the public Internet and they don't belong there, by definition. Where they are seen is at the junction between the Internet and someone's private network -- an end point. They are appropriate there, because end points are where TCP flags should be generated. The middle of the network should concern itself only with the IP header. THE DEVICES IN THE MIDDLE OF THE INTERNET SHOULD NOT GENERATE PACKETS THAT IMPERSONATE AN ENDPOINT OR CHANGE AN ENDPOINT'S CONTENT, except for the very limited allowances covered by the Internet Standards.

EIGHTH: "Network architect Richard Bennett who works for a Router maker explained to me that TCP resets are the standard mechanism used by consumer routers to deal with NAT table overflow, which itself is typically caused by excessive P2P connections." Since this is so standard, I suppose you or Richard can name several popular models available on the market today so that people can test for themselves. (The issue with NAT table overflow seems to be more with the UDP mesh than the actual P2P transport. Because its UDP, a TCP flag like RST is not available. Besides, wouldn't one want to purge older connectionless UDP entries first? (I haven't studied this, but I have seen references to a rational way to prioritize what to trim in a NAT list before overflow occurs.). An ICMP response of Type 3, Code (something) might be the better defense for router manufacturers to use. If the exchange is SYN, RST, then the application developer decides whether or not to retry immediately. If the exchange is SYN, (no response), then the TCP protocol will retry a few times at increasing intervals until some app or protocol-configured timeout. The second way is probably the better way to go if predictability of operation is desired.

NINTH: "Are we to believe a single software tester or three networking experts?" George, congratulations on

working in networking for a Fortune 100 company. The fact is, I am a networking expert with years of intimate experience with the internal and external networks of a FORTUNE FIFTY company (I frequently participated in the change control board managing said networks). B.F.D.! I hate the word, expert -- it tends to signify that "I've arrived" when the fact is that I'm always learning. Expertise is a journey. That said, I'm sure you could show me things that I don't know, and I could do the same. That "professional curiosity" and openness to ideas and change is a trait that you often seem to lack. I was shocked when you said the things I quoted in my fifth point, above, because it refutes a position you previously held strongly. You've understood what I've been saying all along and somehow learned it and adopted it. This is not like you. You usually wear "expert" like a badge. YOU have "arrived." YOUR statements stand above all. It's a small weakness of yours, George. You are an expert but without evaluating alternate understandings and views, you become a parrot. Besides, I'm not a single software tester. My tests were repeated and validated by two separate entities, not to mention the experts on the first Harvard panel, one who borrowed a neighbor's connection to repeat my tests.

TENTH: "The second key problem with Topolski's testimony is that to my knowledge, he has never provided any forensic data from Comcast in the form of packet captures that can be independently analyzed." Read your very next sentence! You said, "Even if Topolski did produce packet captures and we assumed that those packet captures are authentic, one man's packet captures wouldn't be a large enough sample to draw any conclusions by any legal or scientific standards." So then why is this a key problem for you? I weighed whether or not to post Wireshark traces, and decided against it because it takes knowledge of how the TCP protocol should work to understand what one is seeing. So I opted to take it up a level, and describe what is happening and what it looks like when it happens. I asked people to repeat it. They did. Had you called me as the Technical Editor at ZDnet, I would have walked you through this, too, which would allow you to capture your own evidence. I did this with several journalists. And since the interference was constant, there was no occasion where I did not successfully demonstrate it. Apparently I did well enough, since my findings were independently validated. Did you read my original posting on DSLReports and my replies that followed? I posted several artifacts and evidence of the blocking along the way. Did you go to the TorrentFreak articles from August (which started the explosion of Blogosphere interest) and see the evidence posted there? Although uncredited, it came from me. George, if the purpose of this paragraph is to promote the idea that what's actually happening is all in my head, then give it up. **EVEN COMCAST ADMITS THE INTERFERENCE.** Nobody should just believe me. But I did post under my name with my own reputation at stake. In doing so, I hoped that people would consider the possibility and test it for themselves. They did.

ELEVENTH: "Furthermore, even if a TCP reset was used by Comcast at 1:45AM, we cannot assume that there was no spike in congestion at 1:45AM." That's true, but doesn't consider the fact that since March of 2007, I've never found any period where the interference was not evident. I had the same files up for Gnutella, 24/7, from March to May of 2007, and EVERY SINGLE REQUEST to obtain those files resulted in Comcast tearing down the connection with a reset.

George, call me sometime.

By Robb Topolski on 4/28/2008 12:34 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

"I noticed that you consulted with Brett Glass and Richard Bennett before posting your article. You certainly have my permission to consult with me, too."

Because I was only seeking expert and accurate advice on networking.

"and because P2P uploads is what Comcast is tearing down. With BitTorrent, only 3-4 streams are allowed to upload simultaneously. This is a major flaw in YOUR testimony to the FCC and other bodies (I'm referring to your article and graphic showing a single connection comparison to 11 connections).

Once again you show that you do not have a firm grip on the facts. BitTorrent DEFAULTS to a SOFT limit of 4

streams PER TORRENT. There's no hard rule that it must be 4 streams and if you're using less than 80% of your upstream target, it will go over 4 streams PER TORRENT. There is no limit to how many Torrents you can have simultaneously. It is quite common to have 3 torrents running with 12 upstream TCP sessions.

Furthermore, even if all YOUR upstreams are blocked by TCP resets issued by Comcast, that Torrent as a whole is still alive UNLESS you're the only seeder which is usually not the case for a normal Torrent. You're using contrived examples of files that no one wants like a legal copy of the King James TEXT or 19th century Barbershop. These files are more appropriately distributed using HTTP or FTP. It was no accident that I couldn't find a single Torrent for these two examples of yours on Mininova.

"THIRD: "The second problem with Jon Peha's testimony is his implication that the user must take some sort of action to resume the BitTorrent connection or else the connection wouldn't resume." Jon Peha is entirely correct, as long as you understand "user" as being the application and not the person."

Jon Peha was talking about a human redialing the phone and clearly implied that this was a burden on the USER. The word "user" even in the geek world is defined as a person. Had Peha used the word "client", then that would be more ambiguous.

"FOURTH: 'So even if all TCP streams in a Torrent were simultaneously blocked for a short period of time, it will quickly resume by itself and eventually finish the file transfer. Therefore this is by definition a "delay" and not a "blockage".' YOUR STATEMENT IS NOT TRUE when the affected peer is trying to upload unique content"

It's still technically a delay because if you leave the seed running, it will eventually go through. However, I did say that TCP resets have a harsher effect on rare Torrents. I also did point out that you have a superior FREE alternative using the gigabyte of web space provided by Comcast to distribute your rare content that works 10 to 20 times better than using a single P2P seed. Your argument is contrived.

"SEVENTH: "[Topolski] insists that the TCP reset mechanism isn't common in network devices and declared that I was wrong in my testimony." I stand by my statement. They are not COMMON in end-user network devices, although some exist."

Most people use the NAT boxes that their ISP provides and those boxes regularly use TCP resets to preserve their NAT table space. This is also quite common in Linux or BSD and a lot of "end-use network devices" are based on Linux. That "end-user network device" at the University of Colorado was the real culprit when they tried to blame Comcast for the TCP resets.

"The fact is, I am a networking expert with years of intimate experience with the internal and external networks of a FORTUNE FIFTY company (I frequently participated in the change control board managing said networks)."

You participated in a change control board for a Fortune 50? Gosh, Richard, Brett, and I are in awe and we would like to know when we can start taking networking classes from you. Heck, I only DESIGNED and BUILT the network, security, and server infrastructure by myself.

"They did. Had you called me as the Technical Editor at ZDnet, I would have walked you through this, too, which would allow you to capture your own evidence. I did this with several journalists."

You're assuming that I need you to walk me through Wireshark or that I am using Comcast so let me clarify for you. I do not need you to explain to me how to use a sniffer. I do not use Comcast so it would not be easy for me to test Comcast's network. Tests from a single person are too small a sampling point and Comcast doesn't dispute issuing TCP reset packets. I'm more interested in looking at something like the Vuze data if they can make it more granular and plot out the time of day.

By host [George Ou -rmt-] on 4/28/2008 12:44 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

"Furthermore, even if all YOUR upstreams are blocked by TCP resets issued by Comcast, that Torrent as a whole is still alive UNLESS you're the only seeder which is usually not the case for a normal Torrent."

The *whole* point here is that a "normal torrent" *necessarily* has to start as a "rare torrent" *before* it can become a "normal torrent". EVERY SINGLE TORRENT HAS TO START AS A RARE TORRENT. So if Comcast essentially blocks rare torrents from seeding to other BitTorrent users, then that rare torrent cannot ever become a normal torrent, even if the content of the torrent is legal and something that many people want to download.

"I also did point out that you have a superior FREE alternative using the gigabyte of web space provided by Comcast to distribute your rare content that works 10 to 20 times better than using a single P2P seed. Your argument is contrived."

I'm sorry, but tests from a single person are too small a sampling point to say whether Comcast's "superior FREE alternative" is actually superior or actually an alternative. Besides, that ignores the whole point of using BitTorrent. If you put something up on Comcast's "superior" "alternative", no one is going to be able to download it once you take it down. In contrast, a download from BitTorrent is supported by anyone who is currently downloading or seeding the file. That means that a user can seed a torrent for a day or two, take his copy down, and the file is still available for download via the torrent. Torrents exist precisely as an *actual* alternative to server-client paradigms such as Comcast's "superior" "alternative".

"You participated in a change control board for a Fortune 50? Gosh, Richard, Brett, and I are in awe and we would like to know when we can start taking networking classes from you. Heck, I only DESIGNED and BUILT the network, security, and server infrastructure by myself."

Oh, wait, I get it. *You* are allowed to namedrop and say "I worked at a Fortune 100 company", but when somebody else comes along and does it, you dismiss that out of hand. That's rich.

By Simone Manganelli on 4/28/2008 3:07 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

Nice slams. Far be it for me to argue with an "Expert."

1. Multiple Torrents: Fine, but then the comparison isn't 11:1 it's 4:1 or 7:2 or 10:3 --- something to that degree. 11:1 is just false.

2. Define user: <http://www.atis.org/glossary/definition.aspx?id=336> ... programmers of interfaces use the word user all the time to refer to the application or process that uses them.

3. "temporarily blocked" is all we need to know. That it will eventually go through didn't happen for me on Gnutella over 2-3 months and assumes someone is willing to wait for it to "eventually" happen. Meanwhile, they're (you used the word) "blocked." And Comcast does not block, (or delay, or degrade) remember?

4. I'm still waiting for the specific model numbers of NAT devices that send their own RSTs. Since it is so common, you SHOULD be able to name one (if not several)

5. You took me to task for not providing evidence, then said it wouldn't matter anyway. Still you labeled it a top concern?

By Robb Topolski on 4/28/2008 2:21 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

Robb, it's amazing to me that a man who claims to be a networking expert isn't aware of the fact that NAT breaks the so-called end-to-end model. A NAT transforms IP addresses and ports, and typically alters sequence numbers as well. All the commercial NATs will send TCP RSTs when their table space is exhausted, and the typical cause of that is P2P probing for faster connections.

By Richard Bennett on 4/28/2008 2:53 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

"Nice slams. Far be it for me to argue with an "Expert." "

Robb, you have demonstrated over and over again that you do not grasp even the most fundamental facts. You have shown me nothing to indicate any level of expertise in these matters of networking or BitTorrent. Sitting on a few change control committees does NOT make you a networking expert, building and designing networks for large scale enterprises in my example or designing protocols, products, applications like Bennett or Glass does.

"1. Multiple Torrents: Fine, but then the comparison isn't 11:1 it's 4:1 or 7:2 or 10:3 --- something to that degree. 11:1 is just false."

First of all Robb, you are dead wrong that BitTorrent restricts you to "ONLY 3 or 4" upstream TCP flows and even the default setting has a soft cap of 4 flows. It's interesting that you quickly forget what you said just one post ago.

Second, the context of the 11:1 chart is talking about a P2P user versus a normal user. A normal user typically doesn't even have 1 active session open so 11:1 is being quite conservative. If you factored in the persistence factor, it's more like 100:1 in terms of how much bandwidth a P2P user hogs. You insist on talking about upstream TCP sessions on a single torrent but you're conflating the issue with my bandwidth hogging charts.

Third, we're talking about the health of a P2P Torrent or swarm so the total number of TCP streams across all the seeders is what is relevant and not the total number of stream sessions per individual seeder.

3. "temporarily blocked" is all we need to know. That it will eventually go through didn't happen for me on Gnutella over 2-3 months and assumes someone is willing to wait for it to "eventually" happen. Meanwhile, they're (you used the word) "blocked." And Comcast does not block, (or delay, or degrade) remember?

A temporary block is by definition a delay. It's like sitting in traffic at a stop sign or light where you are temporarily blocked from crossing the street. Eventually you'll be permitted to flow. Gnutella or P2P in general isn't the right protocol to be transferring rare files. You could have placed those on your FREE Comcast-provided web space and it would have worked much better.

Now here's one place where you and I could probably agree on. I believe that Comcast should provide their P2P users (in addition or in place of their 1 GB HTTP web space) a 384 Kbps upstream-capped P2P seed server in their datacenter where you control what's being seeded so long as it's legal content. That would not only alleviate your upstream pipe, it would alleviate upstream congestion on the DOCSIS 1.1 network. Then you get to effectively seed all you like without making an impact on your upstream or your neighbor's upstream. Instead of you screaming for no network management which results in massive network congestion, why don't you join me in asking for a sensible solution?

"4. I'm still waiting for the specific model numbers of NAT devices that send their own RSTs. Since it is so

common, you SHOULD be able to name one (if not several)"

I already named examples in my last post. I pointed out Linux, BSD, and Netgear which was used by the University of Colorado when they tried to blame Comcast. Richard just set you straight on this topic as well that ALL commercial NATs (included in all commercial routers) will issue TCP resets when their table space is exhausted.

"5. You took me to task for not providing evidence, then said it wouldn't matter anyway. Still you labeled it a top concern?"

I merely pointed out SEVERAL problems with your argument. Logic should tell you that the second argument doesn't negate the first one.

By host [George Ou -rmt-] on 4/28/2008 2:57 PM

Simone, do you not understand that 6 megabits is 15.6 times faster than 384 Kbps?

Simone says: "I also did point out that you have a superior FREE alternative using the gigabyte of web space provided by Comcast to distribute your rare content that works 10 to 20 times better than using a single P2P seed. Your argument is contrived."

I'm sorry, but tests from a single person are too small a sampling point to say whether Comcast's "superior FREE alternative" is actually superior or actually an alternative."

Simone, do you not understand that 6 megabits is 15.6 times faster than 384 Kbps? A P2P seeder is by definition capped to 384 Kbps if that's their top upload speed. Any server in a data center can easily burst to 100 Mbps. This is just common sense.

"Oh, wait, I get it. *You* are allowed to namedrop and say "I worked at a Fortune 100 company", but when somebody else comes along and does it, you dismiss that out of hand. That's rich."

Simone, do you not understand the difference between a professional Network Engineer who designs and builds networks to someone who merely sat in on a few network infrastructure meetings?

By host [George Ou -rmt-] on 4/28/2008 4:03 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

@ "You participated in a change control board for a Fortune 50? Gosh, Richard, Brett, and I are in awe and we would like to know when we can start taking networking classes from you. Heck, I only DESIGNED and BUILT the network, security, and server infrastructure by myself."

George, loose the attitude there 'tough guy'. It's not necessary and it degrades the debate besides who knows Robb might be closer to the truth that most realize.

By humble Importance on 4/28/2008 4:25 PM

Humble, Richard and I have debunked everything Robb said here

"George, loose the attitude there 'tough guy'. It's not necessary and it degrades the debate besides who knows Robb might be closer to the truth that most realize."

Humble, Richard and I have debunked everything Robb said here point-by-point. I wouldn't normally go around saying my credentials are better than someone else's but when that person goes before the FCC and in public and says I am wrong when he is the one who does not know what he is talking about, then I will stand up and defend myself.

Topolski not only made incorrect statements at the Stanford hearings; he again made incorrect statements here and in a letter to the FCC docket. In both cases he has been fully debunked. When Robb Topolski who has no network engineering credentials come out with bad information and claims he is an expert, I am going to call it like I see it.

Now if you have any technical points you would like to rebut, please state them here. This isn't the place to be crying about someone being mean to someone else.

By host [George Ou -rmt-] on 4/28/2008 4:39 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

If you want to see a great example of somebody touting his alleged credentials, see Topolski's FCC filing where he claims to know more about networking than the entire Comcast corporation. It's truly rich:
http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6519870563

By RichardBennett on 4/28/2008 4:47 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

George, this contest of expertise is all beside the point. My expertise is what it is. I'm not claiming superiority, and I find your responses both childish and unprofessional.

The facts are what matter to me.

Best wishes,

Robb

By Robb Topolski on 4/28/2008 5:00 PM

In reply to Richard Bennett on Robb's letter to FCC

Robb explains to Comcast that he can teach them something:
"I wish to help you as well as everyone else accurately understand these issues ahead of the upcoming meeting at Stanford. You (Mr. Cohen) are not a technologist and you can only report what you've been told."

Robb's bio:

"My bio is file with the Free Press's original filings in this matter. In short, I'm a recognized testing and networking professional with experience spanning over 25 years. In 2004, I earned my qualification as a Certified Software Quality Engineer by the American Society for Quality. In 2006, I was awarded the Most Valued Professional status in the Networking area by Microsoft. Indeed, my professional career in Software Validation and Quality Assurance is weighted heavily in networking."

I still don't see where he designed or built any Ethernet or IP networks, designed any protocols, built any

products, wrote any software, or showed any indication that he is anything more than a software tester yet he feels that he is more knowledgeable than George Ou, Richard Bennett, Brett Glass, and David Cohen.

The bottom line: Everything Robb Topolski has said to the FCC and everything he has posted here has been filled with errors and thoroughly debunked here.

By host [George Ou -rmt-] on 4/28/2008 5:35 PM

Robb, every "fact" you have raised have been debunked

"George, this contest of expertise is all beside the point. My expertise is what it is. I'm not claiming superiority, and I find your responses both childish and unprofessional."

Robb, when you call a group of actual networking experts wrong before the FCC and before the public, and you arrogantly tell a Comcast officer you are there to educate him, you better be prepared to back up your arguments and your credentials.

"The facts are what matter to me."

You say you care about the facts but every "fact" you have raised have been debunked and you can't rebut a single point. The fact that you are now reduced to crying foul about people questioning your credentials pretty much confirms that you don't have a shred of credibility.

By host [George Ou -rmt-] on 4/28/2008 5:11 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

@ "Now if you have any technical points you would like to rebut, please state them here. This isn't the place to be crying about someone being mean to someone else."

We are not crying we are just stating if this was face to face we wouldn't cry either, we'd just kick you ass till you stopped.

Unfortunately, we are probably more aware of the details than anyone on this list (technical and otherwise) but for reasons that can't be said we can't comment.

By humble Importance on 4/28/2008 5:35 PM

So faced with the facts, Humble Importance wants to threaten violence.

"@ 'Now if you have any technical points you would like to rebut, please state them here. This isn't the place to be crying about someone being mean to someone else.'"

We are not crying we are just stating if this was face to face we wouldn't cry either, we'd just kick you ass till you stopped."

So faced with having to answer some hard factual questions, Humble Importance wants to threaten me violence. That's pretty typical. Now were you also the guy at the Stanford Hearings that threatened violence on George Ford during public comments? I wouldn't be surprised if you were.

"Unfortunately, we are probably more aware of the details than anyone on this list (technical and otherwise) but for reasons that can't be said we can't comment."

Sure, we'll just take your word for it. You sound about as credible as Robb Topolski. Really, you do!

By host [George Ou -rmt-] on 4/28/2008 5:53 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

Guys, this is pointless.

Robb is merely trying to milk the 15 minutes of fame he got when he noticed the RST packets from the Sandvine appliance. (In fact, they are nothing unusual. Any knowledgeable network engineer would already know that these appliances emit RSTs -- as do other network management devices, such as WebSense. In fact, the technique is mentioned in a patent).

Robb's "discovery" was trumped up by Free Press, Vuze, "Save the Internet," and other activists looking to bash Comcast, and he suddenly found himself in the spotlight... even though those of us who are familiar with firewalls, COPA compliance, and P2P mitigation just said, "Duh! We know they do this. So what?"

In any event, to stay in the spotlight, Robb must continue to trumpet this "great" discovery as if he'd discovered a new continent -- and proclaim that the technique constitutes awful skullduggery on the part of e-vile corporations rather than simply a good way of cutting off a TCP session.

I really wish there were some way simply to satisfy Robb that yes, he's gotten enough attention from pointing out what truly informed people already knew... and that he can go home already. It's getting tedious to see him popping up in forum after forum with no new information, repeating the same misleading accusations of "forgery."

Because my fellow WISPs mostly survive by working hard and keeping a low profile (I've gotten lots of e-mails from them telling me they're too busy to help me debunk Robb's accusations), I'm out there nearly alone defending my industry. (A few other WISPs have at last agreed to stop forward -- for which I cannot thank them enough.) In any event, I've been forced to devote a \$10,000 budget to travel, speaking, lobbying, and publishing just to set the record straight now that "SaveTheInternet" has been trumpeting Robb's misleading assertions. That's money that won't go into expanding our rural network to unserved locations this year. Thanks a lot, Robb; you've done a lot to help people get fast, secure access to the Internet.

Not.

By Brett Glass on 4/28/2008 5:59 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

Dear Brett, George, and Richard:

You guys have debunked nothing. Simply claiming that you have "debunked" something without providing any evidence or artifacts that would be contrary, is not debunking. For example, countering something I said with a statement that "all routers issue resets" proves nothing at all.

My 15 minutes should have been over in May of 2007. Instead, Comcast took half a year to turn from an outright denial to weasel-worded admission. Like so many events in life, it isn't the bad deed that causes the stir, it's the attempt to cover it up that really causes problems.

Brett, I can't "trump up" the truth. The Associated Press covered the story because it was news, nor did anyone "trump up" their own independent confirmation. Comcast implemented this secretly to avoid scrutiny, not

because every agreed "Duh! We know they do this. So what?" The FCC held two hearings (with more to come, probably) because this behavior is so entirely questionable.

George, the Comcast officer told the FCC that he was not an engineer and did not have first-hand knowledge of what was really happening. When he reiterated that canned paragraph, I stewed on it for a week while it sat unchallenged before I finally decided that it demanded a response and that Cohen deserved the benefit of the doubt. Perhaps, he really did not know.

Richard, my bio is simply my bio. I summed it up at the top of this thread in my SIXTH point. Based on my skills and experience to the best of my ability, I continue to stand by every fact that I have posted. When I'm wrong, I have the decency to readily admit it.

That's it. If your mission today was to waste a lot of my time, rest peacefully knowing that you've accomplished your goal.

By Robb Topolski on 4/28/2008 10:58 PM

Topolski, your posts here and your testimony have been filled with inaccuracies

Debunked Topolski claim #1: TCP resets aren't common in networking devices and "end-user" routers.
FALSE: Richard Bennett explained that TCP resets are the STANDARD method used by NAT to avoid overflows. Since every router has a NAT table, every standard router has implemented TCP reset. So the question isn't which routers use TCP resets; the question should actually be which routers DON'T implement TCP resets to close off connections.

Debunked Topolski claim #2: BitTorrent is restricted to ONLY 3 or 4 upstream TCP sessions.
FALSE: BitTorrent defaults to 4 upstream sessions soft cap per torrent which can be exceeded by default and it can be user reconfigured. Multiple torrents mean even more upstream sessions which give BitTorrent and other P2P applications a huge advantage on a congested network over non P2P users. This explains why 10% of Japan's population using P2P can account for 75% of all Internet traffic.

Debunked Topolski claim #3: Comcast conclusively issues TCP resets when the network isn't congested.
FALSE: Topolski cannot make any such claim. No forensic data was submitted. Even if it were submitted, it's not a large enough data sample. Even if the resets came from Comcast, Topolski can't know whether the network is congested at 1:45AM in the morning.

Debunked Jon Peha claim #1: A TCP reset TCP session constitutes blockage.
FALSE: P2P utilizes tens of TCP streams and a few temporarily blocked streams constitute a delay.

Debunked Jon Peha claim #2: Users must take an action and manually redial to reestablish P2P connection
FALSE: P2P clients like BitTorrent will automatically resume where they left off without user interaction. Topolski's attempt to redefine "user" as an application is feeble.

By host [George Ou -rmt-] on 4/28/2008 11:14 PM

Re: Comments on inaccurate testimony at the FCC Stanford hearing

Ou: "Simone, do you not understand that 6 megabits is 15.6 times faster than 384 Kbps? A P2P seeder is by definition capped to 384 Kbps if that's their top upload speed. Any server in a data center can easily burst to 100 Mbps. This is just common sense."

Ou, do you not understand that 16 peers on a single torrent is a low figure for "normal torrents"? Healthy torrents can have well upwards of 3000 peers. So whether or not a single P2P seeder is capped at a 384 Kbps upload

speed is irrelevant, because the "aggregate upload bandwidth" by all peers of a single torrent can easily saturate the download bandwidth of a single peer.

Besides, would you like me to quote you from the FCC hearing at Stanford? I was there. You said that the speeds advertised by ISPs like Comcast are "peak speeds" not average speeds. This means that it does not even require 16 peers to saturate the download bandwidth of someone on the Comcast network. At home, we subscribe to Comcast's internet service with an advertised download bandwidth. In my experience, I have not seen download speeds in excess of 300 KB/s or so, which would translate to a 2.4 Mbps average connection. That means that it'll probably take only 7 peers on a single torrent to saturate my download bandwidth. If you have any experience with torrents, you'll know that 7 peers on one torrent is an extremely *low* number.

This *still* does *not* address the fundamental point that has been made and that you have so far refused to address. *Every* *single* *torrent* starts out as a "rare torrent". If Comcast essentially blocks "rare torrents" from being able to seed to other peers, then these "rare torrents" cannot *ever* become "normal torrents". So in carrying out its "network management policies", Comcast effectively blocks any new content from its users to be provided via BitTorrent.

Ou: "Simone, do you not understand the difference between a professional Network Engineer who designs and builds networks to someone who merely sat in on a few network infrastructure meetings?"

So basically you're saying that anyone who doesn't have your experience in networking cannot possibly bring anything to the table regarding this issue. Users, hackers who like to tinker with technology, open-source developers, even FCC commissioners who cannot hope to understand the depths of TCP networking in a few hearings -- none of these people have the resume to satisfy your demands of what constitutes someone informed enough to comment on these issues.

This is practically the poster-child for an "elitist". It's funny that you tout your credentials, yet the history of the computer and of the internet is littered with people who were just hackers working out of their garage who had no formal experience in the field whatsoever, who were interested in something and gained knowledge and learned something simply by tinkering with it.

By Simone Manganelli on 4/29/2008 12:28 AM

Rare seeders become healthy seeds if they're popular no matter what

"This *still* does *not* address the fundamental point that has been made and that you have so far refused to address. *Every* *single* *torrent* starts out as a "rare torrent". If Comcast essentially blocks "rare torrents" from being able to seed to other peers, then these "rare torrents" cannot *ever* become "normal torrents". So in carrying out its "network management policies", Comcast effectively blocks any new content from its users to be provided via BitTorrent."

Rare seeders become healthy seeds if they're popular even if the initial seeder is reset a few times; it just takes a little longer to light the fire so to speak. Once the torrent becomes more distributed, it is less affected by TCP resets. Furthermore, it's quite possible to get a complete download from a torrent with zero seeders. That's because various people will have various pieces of the puzzle but not necessarily all the pieces. You put them together and it produces a complete file. Comcast current system doesn't go after the pre-seed peers so just because the only seeder gets temporarily blocked doesn't mean the torrent dies since you're still pulling data from other partially completed peers.

"Ou: 'Simone, do you not understand the difference between a professional Network Engineer who designs and builds networks to someone who merely sat in on a few network infrastructure meetings?'

So basically you're saying that anyone who doesn't have your experience in networking cannot possibly bring anything to the table regarding this issue."

I'm saying that someone who arrogantly and incorrectly claims the real networking experts are wrong deserves to be called out and deserve to have their credentials scrutinized. All the assertions that Topolski raised here and elsewhere along with the assertions made by Jon Peha have been debunked point by point here. Nothing more needs to be said or explained.

By host [George Ou -rmt-] on 4/29/2008 12:46 AM
