

A New Approach to Solving the Net Neutrality Conundrum

Steven Shepard

With the recent decision by the Superior Court to vacate the FCC's Comcast decision, the question of how to manage Net Neutrality arguments has become even fuzzier than it was when we thought we had an answer to it. For those of you who have not followed the Net Neutrality drama or who are not familiar with what it is, let me give a brief explanation (those already familiar with the issue can skip ahead).

The legacy telecom networks, which include both cable and telco infrastructures, were designed and engineered for the traffic requirements of voice and rudimentary data - e-mail, for example. Pricing and usage plans were based on the known performance parameters of the network and its users, a model that worked well.

With the arrival of media-rich content, however, that model began to show its limitations. While networks were fully capable of handling voice and data traffic, they began to stagger under the load when iTunes, YouTube, Facebook, Twitter, Bit Torrent, mobile TV, and other applications began to pummel the network with increasing volumes of traffic. At the same time, customers began to demand increasingly stringent levels of not only quality-of-service, but quality-of-experience as well. The network, they maintained, should handle all traffic, regardless of volume, type or time-of-day. Service providers, of course, protested, but customers pointed to the all-you-can-eat service plans they had purchased, backing the service providers into a legal corner that they could not easily extricate themselves from.

Things came to a legal head when cable provider Comcast took steps to protect itself (and its cost structure) from the media onslaught. Using discrete packet analysis tools, they concluded that the biggest offenders were those customers using Bit Torrent (a peer-to-peer file sharing service) to download large movie and software files. Claiming the right to do so under the auspices of network management, they began to selectively block all Bit Torrent packets, thus eliminating congestion caused by the enormous packet flows caused by the service. Of course, there was another driver behind their decision: if they were to allow customers the unlimited right to use these file sharing services, they would soon find themselves in the uncomfortable position of having to invest in added infrastructure - but without any added revenue to cover the cost of doing so.

Comcast made one serious error when it chose to block Bit Torrent traffic: *it told people it was doing so*. Before long, customer advocacy groups became aware of Comcast's actions, and inevitably, so did the FCC. Citing Net Neutrality rules, which, among other things, state that a packet is a packet is a packet, the commission ruled that Comcast had no right to selectively block Bit Torrent and ordered them (and anyone else

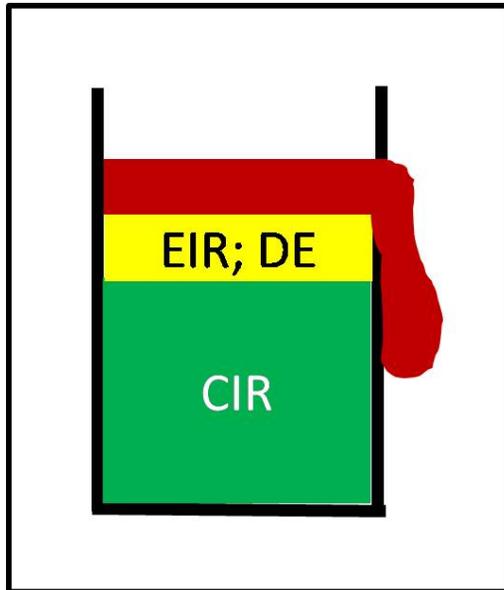
contemplating such action) to cease and desist. This decision seemed to put the Net Neutrality debate to bed, once and for all. The recent decision to rescind the FCC's order, however, brought the whole issue back into the spotlight, and gave the FCC a few options for regaining the authority they believe they should be able to exercise in the interest of protecting network users, often from themselves.

One concern that many advocacy groups have is a valid one. If service providers now have the right to selectively block traffic any time they like under the guise of network management, what's to stop them from blocking traffic that comes in from a competitor's network? For example, what's to stop AT&T from blocking all Skype or Google Voice traffic to force customers to use only their service? At the risk of sounding like a conspiracy theorist, the question is a valid one and demands that the government's various oversight organizations take proactive steps to deal with this issue - and soon.

So let me get back to the matter at hand. It occurred to me yesterday that an answer to the coming Net Neutrality storm may be to look back in time to the days of - wait for it - frame relay. Frame relay may have been a simple and in some ways primitive service, but there was nothing primitive about the way it managed traffic volumes - or customer expectations. Frame relay, I believe, may hold the secret to fixing the Net Neutrality debate.

Frame relay is a switched broadband service that emerged in the early 90s as a solution to the high cost and low flexibility inherent in private line solutions. Working with service providers, enterprise customers would add up the bandwidth they required at each interface, and the service provider would then provision an access facility that would meet that requirement. However, it was actually much more elegant than that. Here's an example. A small-business customer wants to connect a videoconferencing system, five telephones, and five PCs to their router, which will in turn connect to the frame relay network. The video unit requires 384 Kbps to operate properly, while each phone requires 64 Kbps and each PC requires 256 Kbps. The aggregate bandwidth needed, therefore, is 1.984 Mbps, which is effectively a pair of T-1s. That's a costly proposition for a small business.

This is where the elegance of frame relay really shines through. During the discussions with the service provider, the customer realizes that while there are five PCs and five phones in the office, it is rare that there are more than two people in the office at any point in time. Furthermore, the videoconferencing unit is only used sporadically. As a consequence, the most bandwidth that will ever be used in the office, with a few exceptions, is 640 Kbps or so. So the service provider suggests a fractional T-1 of 768 Kbps, thus dramatically reducing the monthly cost for the customer.



The story gets better, however. When the contract arrives for signature, it has several service elements in it that are quite unique. The first is what is called the Committed Information Rate (CIR). The CIR is the amount of traffic that the service provider guarantees it will make available to the customer *at any point in time*. Consequently, the CIR is not measured in bits-per-second; it is measured as *bits in one second*. There is a subtle but significant difference between the two. Bits-per-second is a measure of sustained bit rate, while the CIR (bits in one second) is a measure of the peak allowable traffic load at any point in time. In essence, it's an average. Here's how it manifests itself to the customer in terms of what

frame relay does for them. "At any point in time, we absolutely guarantee that we will make 768 Kbps of bandwidth available to you," says the service provider. "We know that you aren't going to be blasting that much traffic into the network 24 hours a day, so we're actually going to grant you a certain amount of leeway above the CIR. Assuming the network isn't too heavily congested, we'll even allow you to occasionally burst well above the CIR, up to the maximum bandwidth of the provisioned circuit (A fractional 768 kbps circuit, for example, is typically provisioned over a physical T-1). "Don't make a habit of it, but that's sort of our gift to you. Now be aware, anything that bursts above that level (known as the Excess Information Rate, or EIR), will get the Discard Eligibility (DE) bit set, meaning that if the network suddenly gets congested, we reserve the right to throw away that excess traffic to preserve the integrity of our contractual commitments. However, as long as you don't make a habit of abusing our network, we won't make a habit of abusing your traffic. There is, however, a limit to our generosity. While we will allow you to burst above the CIR, we have to draw the line somewhere, and if you go above it, we will automatically, and without notification, discard all traffic above that pre-determined high-water mark to protect the network and all its users from service-affecting congestion. We can't allow the actions of the few to adversely affect the many."

I believe that this is one of the more elegant service management schemes that modern networks have ever seen. I also believe that it may be the answer, in one form or another, to the Network Neutrality debate.

All of us, in fact, the vast majority of us, occasionally abuse the network. We all watch the occasional YouTube video, download movies from various sources, and drain our e-mails, software updates and iTunes selections into our laptops. That behavior, by most measures, could be considered reasonable and acceptable under whatever fair use rules govern networks today. A small percentage of us, however, and it is a small percentage,

do abuse the network by downloading terabytes of movie files, pirated software and streaming audio to our PCs around-the-clock. And while those actions are legally protected under the terms of the service plan, they can cause problems for other customers, especially on shared network segments such as those we find in some cable architectures. Can we deny their right to download large files? Not really; however, we can put into place a pricing mechanism that will allow the service provider to generate enough incremental revenue to cover the cost of network buildouts that are required to support the burgeoning demands of increasingly high-usage, media-rich applications. The answer? Frame relay. A minor adjustment to the existing service plan goes a long way toward addressing the problem.

Consider the following scenario. A “typical” customer (you decide what that means, since it’s something of a moving target) signs up for a service plan. The service plan provisions enough access bandwidth to guarantee that the customer will be able to access the services they want, even taking into account the occasional bursts that result from Friday night’s Netflix download and a healthy serving of the best of YouTube on a regular basis. For the agreed-upon monthly fee, the service provider even guarantees an excess information rate, with the clear understanding that to preserve network integrity there may be times when the service divider has to TRANSPARENTLY (as in ‘we’ll know when they’re doing it’) deny the excess rate. And while the “EIR” is somewhat guaranteed, there is a ceiling on it, and anything above that rate will be subject to a premium charging rate. The premium charge would be set at a level that would (1) not be onerous, and (2) generate enough additional revenue across the vast sea of customers to offset the incremental cost of network augmentation require by the increased levels of customer demand.

This model could even be taken in a slightly different and more qualitative direction. Thanks to a relatively new technology known as Deep Packet Inspection (DPI), the packets traveling across a network can be classified in real-time according to the type of information they contain, thus giving service providers even higher levels of network management latitude than they have ever had before. And while this capability could be used abusively by service providers, it could also be used as a granular management tool, giving them the ability to work with customers to establish ground rules for traffic discard eligibility based on application preferences, content priority or cost. Customers would therefore have say in how they are billed and how their traffic is treated. This capability would eliminate the need or desire on the part of the service provider to use the scheme as a weapon; it would have far more value as a facilitator of proper behavior and as an incentive to customers to engage in responsible network management practices.

The battle between service providers who must walk a fine line between protecting their customers and protecting the network, and the so-called “over-the-top” providers who believe that all packets are created equal, will not be won on the merits of each

side's arguments. The fact is that both arguments have merit, and a loss on either side results ultimately in a loss for the market. Unless the market wins, unless the market gets what it wants, there are only losers. It is critical, therefore, that the decision-makers who have control over this issue allow their pre-conceived notions of right and wrong to fall by the wayside, and instead rely on a new rule. It's quite simple, really: the only way to gain influence in the market today is to give up control of that market. By allowing the customer to have some control over how they engage with and use network resources, they become more responsible for their own destiny and will take more responsibility for scarce resource management. By managing scarce network resources properly, the abundance of digital media content becomes more available to all.

Steven Shepard is the president and founder of the Shepard Communications Group, a Vermont-based firm that provides industry analysis, technical education, management consulting, and media development to technology companies throughout the world. He can be reached at +1-802-878-0486, or via e-mail at Steve@ShepardComm.com. He is also the co-founder of the Executive Crash Course Company, a leading developer of Web-based, multimedia training curricula.