

**Written statement of Dr. David D. Clark**  
MIT Computer Science and AI Lab

FCC public hearing on network management  
February 25, 2007

The issues that the FCC are considering at this hearing are very important, and will materially influence the future of the Internet, as it is experienced by the citizens of the United States. I am pleased that you are considering these issues, and am pleased to provide these observations.

First, by way of context, the issues you are considering here are specifically driven by the emergence of television (and the video experience generally) on the Internet. There are two reasons why video is particularly important and distinctive. First, video really does increase the traffic on the Internet compared to other sorts of content, and traffic is not cost-free to carry. Second, the business model of television (cable television, in particular), which is based on the retailing of content bundles to the consumer as a way to cover the cost of the access facilities, is at odds with the business model of the Internet, where the facilities providers do not sell content as a means to pay for their facilities. It is the collision of these two models, both accepted in their own spheres, that we must manage.

Industry needs to manage the collision. That is not the job of the FCC: it can define behavior that is not acceptable, but it cannot and does not give business advice. But I would plead with all sides here (facilities providers, content providers, advertisers, etc.) to look at the other actors and say “partner”, not “enemy”. ISPs are in a very good position to create products and services that can reduce the costs of delivering TV across an open Internet. I urge them to go down this path and make video a win/win, rather than looking at the content provider (and the consumer of that content) as a “problem” to be managed.

It is reasonable for ISPs to manage the traffic load on their network, and video downloads will indeed materially change traffic levels. But there are several ways to manage congestion and capacity allocation, with different implications. I find blocking of specific applications troubling and regrettable. It raises fears of blocking out of ISP self-interest, but also makes the ISP and the customer look at each other and say “enemy”. The ISP is imposing its value judgment on the customer, and telling the user that his preferred uses of the network are not acceptable. This approach, from the perspective of the Internet, is the whole wrong idea.

At a minimum, full disclosure of what is being done should be required. Disclosure will shine a light on the possibility that the ISP is blocking to favor certain applications. But this is not the only reason for disclosure. Random unacknowledged blocking leads to tremendous user frustration. The Internet is hard enough to debug when it breaks, without adding intentional undocumented injection of faults.

The Internet contains mechanisms to deal with congestion, of course. These have been part of the Internet for over 20 years. However, if an ISP has to deal with some application that does not respond to those controls in the expected way, simple capacity allocation is much preferable from a policy perspective over any form of application-specific intervention. Controlling congestion by rate limiting each customer, independent of application, is value neutral. If there are indeed such applications (or classes of applications) that do not respond properly to network congestion, it should be possible to get some industry consensus on that fact, perhaps in a forum where network operators meet, such as NANOG, the North American Network Operators Group. Having an open discussion in such a context would lend credence to the assertion by an operator that some intervention above and beyond the default congestion controls of the Internet are required.

The best way to deal with high demand is to identify the customers that want to do more and find a reasonable way to price the service so that they become a partner, not an enemy. I believe that this can be done over time, but it will require consumer education and market studies.

I said that video does materially increase the loads on the network, and that usage does cost money. I know that some have the opinion that all the costs of an ISP are fixed costs and that usage is incrementally free, but this is not at all true if the ISP must increase the capacity of its backhaul and interconnection links, etc.

To understand the impact of video, it is important to have a rough model of what usage costs. However, it is very hard to get these numbers, since most ISPs view their cost structure to be confidential. There is a unsubstantiated number, widely circulated (it is to be found on Wikipedia, without citation of source, at [http://en.wikipedia.org/wiki/Broadband\\_Internet\\_access](http://en.wikipedia.org/wiki/Broadband_Internet_access) ) that cost to a broadband access ISP to cover usage is about \$.10 for each Gigabyte per month used by a customer. If we estimate (very roughly) that a typical user, pre-video, uses about 5 GB/month, this would imply that the cost of usage to the ISP is about \$.50/month. For a cable system that charged perhaps \$40/month, less than a dollar goes to usage—the rest goes to fixed costs, including cost recovery on their outside plant, and other costs such as billing, marketing, customer service, and so on.

Using these rough numbers, what might video cost? Assume that a customer watches television for 8 hours a day, and the video stream is 2 mb/s (all these numbers are highly speculative, of course, but I want to offer only an order of magnitude estimate). Those numbers work out to over 200 GB/month, or more that 40 times the pre-video estimate of 5 GB/month. Usage of 40 GB/month would cost the ISP \$20 in usage-related costs.

What can we conclude from these very rough numbers? First, no ISP can absorb an additional \$20/month in usage cost and make money if they continue to charge the same \$40/month. But on the other hand, that \$20 estimate is based on today's costs, and the full TV experience over the Internet is still a few years away. And third, this number does

not take into account any special steps that the ISP (and others) might take to reduce the cost of delivery.

But even if that \$20 drops over time, there is will be real difference in the cost to serve customers that do and do not transfer lots of video. So what sort of pricing can respond to this situation? One idea is different usage tiers for monthly usage. An ISP could charge \$40/month for today's usage pattern, and perhaps \$60 for substantial video downloads. Could the market tolerate that? I think so, and the price and service will only get better if the ISP takes technical and business steps to reduce costs.

As the FCC considers what action to take, and how to balance the interests of the various stake-holders here, I would offer the following considerations.

The collision of the Internet and the video industry, with their two very different price models, will take years to resolve. High-level ground rules will be very important.

I believe that the simple flat-rate "all you can eat" pricing for residential broadband will have to be modified over time to deal with the usage generated by video.

Avoid collateral damage. The open nature of the Web is a crown jewel of the Internet today. We must preserve this open nature, even though video is found on the Web.

The Internet is not neutral, and has not been neutral for a long time. There is lots of discrimination that goes on inside the network: some good, some bad. The vision cannot be a return to a simple world with no discrimination, but to find a line that delineates acceptable and unacceptable.

In particular, discrimination among classes of traffic (e.g. better service for VOIP) makes technical sense and in itself is not a violation of any desired form of neutrality.

What happens outside the public network is equally important to what happens inside. Providers of facilities-based access to the residence often offer both public Internet access and dedicated services over a private network they implement over their infrastructure. I think it is important to consider why building a private parallel net with enhanced service is different from preferentially enhancing some parts of the public Internet.

I believe that the correct approach in this case is not to try to define, ex anti, a bright line between good and bad behavior, but to begin by using some administrative process to begin to define what is clearly acceptable and clearly not acceptable. Over time, the region of uncertainty will be reduced and resolved.