

The Tech Liberation Front

Digg, Network Neutrality, And The Long Tail

May 4, 2007

Over at Cato@Liberty, I've got a post making the slightly obvious point that Digg is a microcosm of the Internet as a whole. Digg, like the Internet as a whole, is an automated and decentralized information-processing system. And just as Digg ultimately faced a choice between allowing the AACCS key to be on the site or shutting the site down, we face the same basic choice as a society: unless we want to shut down the Internet (or radically redesign it, which could amount to the same thing) we've got little choice but to allow some level of illicit content to be traded.

This seems to me to be a nice illustration of a point that I've often tried to make about the network neutrality debate, because it seems to me that the telcos face a similar challenge with regard to their management of their networks. Many of the horror stories proregulatory types tell about a post-neutrality future assume that the telcos have finegrained control over what kind of content flows over their networks. That they're censor liberal blogs, or shut down particular categories of new innovative applications, or sign exclusive deals where (say) one sports website is the official sports website, and all the others are blocked or degraded.

But an ISP attempting to implement such a fine-grained, coercive strategy on a user base numbering in the millions is likely to find their users reacting in creative ways that confound the scheme. Tech-savvy users will immediately start running services on nonstandard ports or tunneling their connections over encrypted links. They'll find ways to camouflage one category of traffic as another, such as making a VoIP session look like a World of Warcraft game. Soon you'd start seeing user-friendly applications available for download to allow moderately tech-savvy users to use the same tricks. And applications developers will start integrating these tricks into their applications, so that the application will automatically detect whose network they're on and use the appropriate countermeasure.

(Geeky aside: it's possible to imagine open source networking libraries that do this automatically and transparently, presenting an API that allows the application developer pretend he's on a normal, open network. Indeed, I bet you'd end up with a situation similar to the situation we saw with open source instant messaging libraries a couple years ago: the telco would introduce new routing polices in an effort to break unauthorized applications. The creators of the circumvention libraries would find a new work-around, publish it, and all the application developers would have to do would be to download the new library and recompile.)

Of course, the telcos could always go for the nuclear option and block all traffic it can't validate as "approved," effectively converting the open network into a closed one But

that would come at a very high price, because there's a long tail of content and long tail of applications. An Internet that only does the things on your ISP's approved list is dramatically less useful than an open Internet, just as Digg would be a dramatically less successful site if it only featured stories that had been pre-vetted by the telco's employees.

So while telcos may have formal control over their pipes, they probably have less practical control over Internet content than is generally assumed. An open network is much more useful to users (and will therefore generate more revenue) than a closed one, but once you have an open network it's very hard to limit how it's used.