Daniel Lyons
@ProfDanielLyons

April 2, 2018 6:00 am | *AEIdeas*

Paid prioritization: Debunking the myth of fast and slow lanes

Technology and Innovation, Telecommunications

Last week, on a panel at the Free State Foundation's always excellent Annual Telecom Policy Conference, Cisco Systems Vice President Jeffrey Campbell highlighted that paid prioritization is "one of the most misunderstood issues" in the telecom policy space. His concerns match a recent op-ed by Sen. Lisa Murkowski (R-AK), which reiterated the importance of the open internet but nonetheless discussed the value of prioritization in capacity-constrained areas such as rural Alaska.

Via Twenty20

Both Murkowski's and Campbell's remarks reflect a growing realization that prioritization can play a positive role in network traffic management. But to understand why, we need to get beyond the "fast lanes, slow lanes" metaphor that has too often dominated the net neutrality debate.

The myth of fast lanes and slow lanes

Net neutrality advocates often argue that without a ban on paid prioritization, internet service providers (ISPs) would divide the network into "a fast lane for those who can afford it — and a slow lane for the rest of us." Playing, perhaps, on the 1990s imagery of the internet as an "information superhighway," this rhetoric envisions broadband networks as segmented into various lanes of travel, with packets sorted into channels that move at different maximum speeds at all times.

But as Campbell explained, this is a myth. All internet traffic on a network moves at the same speed — the speed at which the electrons propagate on the wire. The problem is congestion: what happens when users want to transmit more data than the wire can physically manage at a particular moment. In this case, the network must drop some packets and allow others to go through. The dropped packets then must be resent, which delays the delivery of the service.

The inevitability of congestion and the danger of prioritization bans

Of course, congestion is not constant; it is more likely to occur at times of peak use. The solution to chronic congestion is to expand network capacity. But additional capacity is expensive. It is often uneconomic to build a network with zero congestion at peak time because this would create significant excess capacity at off-peak periods. And a zero-congestion network today may nonetheless face congestion in the future, as consumers' appetites for data grow. So some amount of congestion is inevitable.

So how can we address that congestion? One can drop packets randomly, which seems to align with net neutrality's ethos that all traffic should be treated the same. But there's a problem with this model: Different internet content and applications have different susceptibility to congestion. A user loading an email or a webpage is unlikely to notice if some packets are dropped and resent. But streaming video or FaceTime may buffer, which erodes the consumer's experience and makes the product less reliable. As Campbell notes, an alternative would be to drop packets intelligently, by deprioritizing traffic that is less sensitive and prioritizing traffic that is more sensitive to congestion. This would improve the experience for streaming video (for example) without measurably degrading the web surfer's experience.

See also: States join the net neutrality #resistance: Will it make a difference? (https://www.aei.org/publication/states-join-the-net-neutrality-resistance-will-it-make-a-difference/)

Importantly, Campbell explained that a ban on prioritization is effectively a ban on internet-based services that are sensitive to congestion. Without the ability to secure a minimum quality-of-service guarantee, an edge provider of a congestion-sensitive service is less able to deliver its product to consumers. Murkowski highlights the importance of telemedicine and tele-education applications, both

of which involve congestion-sensitive video transmissions that suffer when all traffic is treated

"equally" at congested nodes. As Murkowski explains, the inability to prioritize congestion-sensitive traffic most adversely affects rural areas, where these video applications are useful and where capacity-constrained networks may experience congestion more often. But Campbell's point is that prioritization bans affect all of us, by deterring innovators from developing new internet-based services that need some protection from congestion.

The role of price

When pressed, some net neutrality advocates will concede the value of intelligent traffic management. The problem isn't prioritization, they claim, but paid prioritization: the protection against congestion in exchange for a fee.

But once one acknowledges the need to prioritize traffic, one then needs a method of prioritization. One solution is a central planning model: An expert (likely either a government bureaucrat or a broadband company engineer) can develop a master list of all internet-based applications and sort it by priority. This raises difficult questions about the sorting rule. Is it based entirely on how quickly the service erodes, or is the expert choosing, say, telemedicine over cat videos because he or she feels telemedicine is more important? This raises the prospect of government or ISPs picking winners and losers, which is precisely what net neutrality is supposed to prevent. The expert may miscalculate an application's sensitivity. And even if the expert gets the list right, it's hard to maintain in a dynamic environment where new services are being added and existing services are being improved, which makes today's congestion-sensitivity calculations less relevant tomorrow.

Alternatively, we can use the price mechanism, which is the way we generally allocate scarce resources (like bandwidth) in a capitalist society. Hayek taught that prices reveal information that markets can use to sort claims on a decentralized basis. An edge provider will only purchase prioritization if its service is congestion-sensitive. When it is willing to do so, and at what price, reveals how susceptible it is compared to other apps. This sorts apps with less error and fewer value judgment than a centrally planned solution.

The concern, of course, is that the price mechanism harms those who cannot afford to pay for prioritization. But these concerns are somewhat overrated. First, apps that are not congestion-sensitive have no need to pay for prioritization. Second, even in a net-neutral world, there are other ways that well-funded companies can — and do — pay to reduce their exposure to congestion, such as using content delivery networks to bypass the public internet. Finally, although players have incentives to use prioritization in an anticompetitive manner, antitrust law protects consumers from such harm, just as it shields us from anticompetitive conduct elsewhere in the American economy.

It's always difficult to understand telecom policy by metaphor, as anyone knows who has tried to teach the *Brand X* case's distinction between Justice Thomas's internet-as-car-dealership analogy and Justice

Scalia's preferred conception of internet-as-pizza-delivery. Too often, metaphors are taken too literally

in ways that harm the policy debate. I share Campbell's hope that we can move beyond the misleading fast lanes/slow lanes metaphor, which hides the significant potential benefits of prioritization to consumers and app developers alike.

Learn more: After the vote: What's next for the Restoring Internet Freedom Order (https://www.aei.org/publication/after-the-vote-whats-next-for-the-restoring-internet-freedom-order/)

Federal Communications Commission (FCC), Innovation, Net neutrality, Regulation, Telecommunications