

advertisement:

THE AMERICAN PROSPECT

The American Prospect's Top Issues:

War on Iraq? Corporate Scandals

online

▶ Archives ▶ Columnists ▶ Issues in Depth ▶ Press Room ▶ Advertising ▶ About Us Search

Volume 11, Issue 10. March 27, 2000 - April 10, 2000.

Print-Friendly Version

Innovation, Regulation, and the Internet

[Lawrence Lessig](#)

In a small hearing room in the House Rayburn Office Building, I met with a group of Capitol Hill staffers to discuss the issue of "open access" in broadband cable. "Broadband" is what policy makers call the next generation of Internet access--faster and always on. Cable is now the dominant mode for serving broadband. And the open-access debate is about whether customers get to choose the Internet service provider (ISP) that serves them broadband cable or must take the ISP of their cable company's choice.

There were maybe 10 staffers in the hearing room, and though kind people call me young, only one was my age (38) or older. They were kids, though they were said to be the ears of the House on this and other issues of cyberspace and its future.

The session started as any law school seminar would. I played the professor--I am a professor--and I laid out an argument about how to understand the present open-access debate. Open access, I argued, has been the rule for narrowband Internet (across telephone lines). Innovation on the Internet is in part due to this rule. We should hesitate before we change that rule, for change may well threaten this innovation.

About five minutes into the session, two staffers came in late. And after about a minute more of my presentation, one of the latecomers had heard enough. Here I was, he objected, arguing that the government should "begin regulating the Internet." Where was the limit? Where would I draw the line? Today I was calling for the regulation of broadband cable; should we also regulate broadband wireless? And if wireless, then satellite too? Was there any stopping this "new" regulation of cyberspace? Was I proposing that we regulate Linux (or "Line-Ucks," as he mispronounced it) because it might become as popular as Windows?

[Query Lawrence Lessig!](#)

The American Prospect Online will be conducting an interview with Lawrence Lessig about the open-source movement. [Post your questions for Lessig on our Online Forums](#), and we will include a selection of them in the interview.

Should Public Policy Support Open-Source Software?

Find out more by reading the [TAP Online Controversy](#).

There is deep confusion about the idea of "regulation" within our political culture and about its relationship to innovation and the Internet. The fashion is to say that regulation harms innovation; that government-backed rules undermine creativity; that the best or most effective policy for regulators is, as Federal Communications Commission (FCC) Chairman William Kennard put it, to allow the "marketplace to find business solutions ... as an alternative to intervention by government."

Any talk about "regulating" cyberspace invites the breathless reply of the impatient young Capitol Hill staffer: Cyberspace was born in the absence of regulation. Don't kill it with regulation now.

This attitude is profoundly mistaken. It betrays an extraordinary ignorance about the history of the Internet, and this ignorance threatens to undermine the innovation that the Internet has made possible. Innovation has always depended upon a certain kind of regulation; the greatest examples of innovation in our recent past evince this reliance. And unless we begin to see the relationship between this type of rule and the innovation it promotes, we are likely to kill the promise of the Internet.

Innovation and Monopoly Power

In 1964 RAND researcher Paul Baran proposed to the Defense Department a plan for a telecommunications network that was very much like the Internet today. It wasn't quite the Internet of today, and it probably wasn't the first such plan. (Baran and Leonard Kleinrock of MIT independently developed this set of ideas in 1961.) But it was, nonetheless, an important and radical change from the existing architecture of telecommunication networks. And the Defense Department took it seriously enough to raise the design with its network experts--AT&T.

AT&T didn't like the plan. At first they claimed the network wouldn't work, but in the end their resistance was about something else. As author John Naughton describes in an extraordinary book, *A Brief History of the Future*, their resistance was about competition. As AT&T executive Jack Osterman put it, "Damned if we are going to allow the creation of a competitor to ourselves."

The AT&T network gave AT&T the power to decide how its network would be used. If an innovator had a different idea about how a telecommunications network should be run, that idea would run on the AT&T network only if AT&T wanted it. AT&T had the power to choose which ideas ran and which ideas did not. It architected this power into its network, and this power was backed up by the force of law. Even if you could, it was illegal to connect devices to the AT&T network that AT&T itself did not approve. The government made sure that AT&T kept control.

This was regulation. It was a power vested in AT&T, both by the architecture of the original network and by government regulations that confirmed the power in that architecture. It was a regime that centralized decisions about how telecommunications should develop. It gave the telephone monopoly the power to protect itself and the opportunity to behave strategically--to decide, for example, not to "allow the creation of a competitor to ourselves."

In this, the regulation of the old AT&T was not very much different from the regulations that still govern cable or broadcasting. There too the law has granted network owners a great deal of power--the power both over the conduit and over the content. Innovation in cable proceeds as cable companies allow just as innovation in broadcasting proceeds as broadcasters allow. The regulation in all three cases buttresses a certain monopoly power over an important part of a communications network.

But at least in the context of telephones, this regulation had an effect on innovation. It stifled innovation. No doubt AT&T spent millions to improve its version of the telephone system. But its version wasn't the only possible one. So long as AT&T kept the keys to the infrastructure, there was little return from thinking differently. Innovators looked elsewhere for projects to develop.

At the core of the open-source and free software movements

lies a kernel of regulation as well. But this regulation is quite different from the regulation that governed AT&T. At its root, open code rests upon a license--upon a kind of law or regulation that controls how this "open code" can be used. Despite the monikers "free" and "open," this license is not forgiving. It is a fairly strict requirement about the uses to which free or open-source software can be put. One does not take open code in the sense one might take a free leaflet from a vendor on the street. A free leaflet one can burn, or box up, or keep from others in a million possible ways. Open code gives the recipient no such power. One takes open code on the condition that one keeps it open--that one distributes it with its source intact, as open as one received it. The open-code movement thus uses law to keep code open. It grants people access to code on the condition that they pass the code along as unencumbered as they received it. (Actually, the licenses are many, and their details different, but this summary will suffice for my purposes here.)

This regulation, like the regulation of the old AT&T, has a consequence for innovation too. But its consequence is quite different. The law in open code means that no actor can gain ultimate control over open-source code. Even the kings can't get ultimate control over the code. For example, if Linus Torvalds, father of the Linux kernel, tried to steer GNU/Linux in a way that others in the community rejected, then others in the community could always have removed the offending part and gone on in a different way. This threat constrains the kings; they can only lead where they know the people will follow. The resource--the source code--is always out there to fuel a revolution, protected by a license from capture by any single person or corporation.

This consequence in turn has an effect on innovators. It assures developers on an open-code platform that the platform cannot behave strategically, that it can't turn against them. If a developer writes a browser for an open-code operating system, there is no way the operating system can force a competing browser off the platform. Even if the browser is bundled inside the operating system, the bundle can always be undone. As the source code is always available, competitors can never be stopped from bundling the operating system differently. This effect marks an important difference between open- and closed-code systems: Whether or not you believe that Microsoft tied its browser to its operating system by linking the code of its browser to the code of the operating system so it could not be removed by competitors, it is clear that an open-source operating system could never be accused of the same charge. There is no way for an open-source operating system to tie itself to any particular path of development. That power is removed by an architecture that ensures that the source is always available. And that architecture, supported by the force of law, guarantees that consumers have jurisdiction over the innovations that will prevail. Thus, unlike the regulation supporting the old AT&T, the regulation in open code operates to decentralize control and to ensure that many have the opportunity to innovate; it guarantees that no single vision of a product gets the power to capture that product. Only the market gets that power.

The Case for Competitive Neutrality

The Internet is the fastest-growing computer network in history. It is not, however, the first computer network. There were many before it, many of which were extremely well-funded. Something, however, was different about the Internet, something in its design.

In the view of many, the critical difference is a design principle that network architects Jerome Saltzer, David P. Reed, and David Clark call "end-to-end." This model regulates where "intelligence" in a network is placed. It counsels that intelligence be placed in the applications. As described by

Saltzer, "end-to-end" says: "Don't force any service, feature, or restriction on the customer; his application knows best what features it needs and whether or not to provide those features itself." Build the network to give the application or users control over the service; don't allow the network any such control. The network is to remain stupid, and intelligence is to reside at the ends.

End-to-end was initially chosen as a technical principle. But it didn't take long before another aspect of end-to-end became obvious: It enforced a kind of competitive neutrality. The network did not discriminate against new applications or content because it was incapable of doing so. The network can't tell the difference between a packet carrying Republican speech and a packet carrying Democratic speech; it doesn't notice the difference between a packet sent from a Windows operating system and one sent by Linux; it can't filter out the streaming of video from the streaming of audio. The network is designed not to know these differences, but simply to take the packet offered and route it as it is addressed. This doesn't mean that users can't discriminate. The point of end-to-end is not that everything goes; it is to locate the power to discriminate in the users--they choose--and to remove that power from the network. The principle thus regulates the power to discriminate. It requires that the network have none.

This regulation too affects innovation. Just like the license governing open code, end-to-end means that the network owner can't pick and choose which applications or content will run. As the network can't discriminate, the test of whether new content or applications run is thus not whether the network owner likes it, but whether the content or application can be coded in an IP protocol. If it can, it will run; and if it is desired, then it will become dominant. Like open code, the principle of end-to-end vests control over the evolution of the Internet in (the many) developers and consumers, and not in (the few) network owners. Like open code, it is a regulation designed to enable innovation.

The consequence of this principle has been profound. By keeping itself open to evolution, the network has developed in ways that no one would have imagined at the start. At each stage, there have been pressures to optimize on the present model, and the commitment to end-to-end has avoided such calcification. As Saltzer, Clark, and Reed note, had the network been optimized in the 1980s for telephony, as many thought it should, the World Wide Web would not have been possible. A commitment to simple and stupid networks has produced an opportunity for surprising and radical innovation.

End-to-end differs from open code, however, in an increasingly important way. Unlike the restrictions that govern open code, the principle of end-to-end is not enforced by law. It is perfectly possible, and in the main, completely legal, to build technologies that violate end-to-end, and then to integrate those technologies into the Internet. Many companies have, and among the technologies being proposed for the future, many more will. Thus, rather than a rule, end-to-end is a norm among network architects. And like many norms, it is increasingly becoming displaced as other players move onto the field.

But there is one part of the Internet where end-to-end is more than just a norm. Here the principle has the force of law, and the network owner cannot favor one kind of content over another or prefer one form of service over another. Instead the network owner must keep its network open for any application or use the customers might demand. Competitors must be allowed to interconnect; consumers must be allowed to try new uses. In this part of the Internet, "open access" is the rule.

This part of the Internet is--ironically enough--the telephone network, where because of increasing regulation imposed by the D.C. Circuit Court of Appeals in the 1970s--leading to a

breakup of AT&T by the Justice Department in 1984 and culminating with the Telecommunications Act of 1996--the old telephone network has been replaced with a new one over which the owner has very little control. Instead, the FCC spends an extraordinary amount of effort making sure the telephone lines remain open to innovators and consumers on terms analogous to the terms required by an end-to-end principle: nondiscrimination and a right to access.

The FCC is convinced that this regulatory burden is severe and costly to maintain. And no doubt it is costly. But the question is not simply how much the regulation costs; it is also about its benefit. What is the benefit of effectively enforcing end-to-end on the telephone system?

In my view, the benefit has been the Internet. Though the Internet proper was initially a network among universities, had it not been for the ability of ordinary consumers to connect to the Internet, that network would have gone nowhere. (Universities are fun, but they aren't enough to fuel commercial revolutions.) Ordinary consumers connected to the Net across phone lines. And had it not been for the open-access rules that the government imposed upon telephones, the telephone companies would most likely have behaved just as every network owner in history has behaved--to control access and use architecture to minimize competition. If it hadn't been as cheap to dial a local bulletin-board system (BBS) as it was to dial a local friend; had the Baby Bells kept the power to force customers to a Baby Bell ISP; had the government not insisted that competitors be connected and had it not policed pricing to ensure nondiscrimination--had it not, in short, used the power of law to force a competitive neutrality onto the telephone system, the telephone system would not have inspired the extraordinary innovation that it did.

By keeping the network neutral, by keeping it open to innovation, the FCC has made possible the extraordinary innovation that the Internet has produced. Open access was the rule; a regulation produced that rule.

Competition Policy and Innovation

There is a lesson to be drawn from these three spaces of innovation, a lesson about the relationship between innovation and the power to control. Open code, end-to-end, and open access all seek the same result: a platform where the right to innovate is protected. To this end, they all use a form of regulation to disable a power to control. Open code uses contract, end-to-end uses norms, and open access uses law. This regulation, while fundamentally different from the regulations that gave us the original AT&T, is still regulation. Its aim is to coerce behavior that we would not expect network owners and coders to choose voluntarily, at least after they've gotten control over the network or over the code.

Now of course my point is not that all control stifles innovation; it is not that corporations inhibit rather than build innovation. How much innovation is protected by these regulations is a hard question. How much less innovation there would be if these principles of the original Net were ignored is also a hard question. We have no good way to measure the effect of these regulations protecting innovation. We have no good way to tell whether in fact they were necessary.

But neither do we know enough to say the opposite. We can't say that open code would have flourished as fully as it has without its strict license, or that the Internet would have grown just as it did without the norm of end-to-end; and we don't know whether the Internet would have flourished without the FCC's control over the Baby Bells.

And yet now we must decide whether the same principles of open access and end-to-end should also govern

broadband--whether it be cable broadband, wireless broadband, or broadband through telephone wires. How should we answer this question in the face of what we don't know? How should we resolve it when we can't be certain?

In my view, our bias should be in favor of what has worked unimaginably well. Having tripped onto this environment of extraordinary innovation, we should be cautious before we allow it to be changed. If we can identify the principles that have distinguished the Internet from earlier, less successful networks, then these principles should guide us in choosing rules to govern networks in the future. End-to-end, enforced through open access, has been a central part of the Internet revolution. At a minimum, the burden should be on those who would compromise that principle to show that it will not take away from the innovation we have seen so far.

The choice is not between regulation and no regulation. The choice is whether we architect the network to give power to network owners to regulate innovation, or whether we architect it to remove that power to regulate. Rules that entrench the right to innovate have done well for us so far. They should not be repealed because of a confusion about "regulation."

[Stop Squinting!](#)
[Subscribe](#) to *The American Prospect*.

At the end of my day on Capitol Hill, I met with a smaller group of Senate staffers to discuss the very same issues of open access. These staffers were different: They were older, they seemed to have been around much longer, and they were much more aware. Now I felt like the kid, and the more I described this ideal of end-to-end and its relationship to open access and the principles of the Internet, the more impatient these staffers became.

A deal would be struck, I was told. At least the major ISPs would get access to the cable network. As Chairman Kennard had said, business would find its "business solutions"; access would be granted on terms set by the network owners. The open access of narrowband Internet would not be possible, but neither would closed access be allowed. A compromise would be found.

This is Washington's version of the Internet. There isn't a problem so long as the big guys can buy access. And while, to their great credit, Steve Case of AOL and Gerald Levin of Time Warner have actually pledged their networks to operate closer to the open-access ideal, I couldn't help but feel that a battle to defend the original principles of the Internet was over. If there is principle, it isn't visible in D.C. If there was innovation to protect, it was only the innovation Hollywood might imagine. Our political culture would in time transform the Internet into the shape of everything else. We might recognize the original Net in what Washington produces. But as Kevin Werbach of Release 1.0 put it, whatever architecture it will produce, it has little to do with what the Internet was.

We have the opportunity to preserve the original principles of the Internet's architecture and the chance to preserve the innovation that those principles made possible. But that opportunity will require a commitment by us, and by government, to defend what has worked and to keep the Net open to change--a regulation to preserve innovation. ☐

[Lawrence Lessig](#)

Copyright © 2000 by *The American Prospect, Inc.* Preferred Citation:
Lawrence Lessig, "Innovation, Regulation, and the Internet," *The American Prospect* vol. 11 no. 10, March 27, 2000 - April 10, 2000. This article may not be resold, reprinted, or redistributed for compensation of any kind without prior written permission from the author. Direct questions about permissions to permissions@prospect.org.

[Archives](#) | [Columnists](#) | [Issues in Depth](#) | [About Us](#) | [Press Room](#) | [Advertising](#) | [Forums](#) | [Home](#)