The Edge of Tomorrow

Author : James Grimmelmann

Date : November 28, 2025

Tejas N. Narechania & Scott Shenker, *How to Save the Internet*, __ **Berkeley Tech. L.J**. __ (forthcoming), available at SSRN (Mar. 18, 2025).

Every time I teach Internet Law, I start by lying to my students about how the Internet works. I tell them the finely crafted story of how routing, packet-switching, and layering combine to produce a profoundly modular, decentralized, and standardized worldwide network. The only problem is that the Internet doesn't work that way anymore, and hasn't for years. Companies like Akamai, Cloudflare, and Amazon operate such massive networking infrastructure that they have warped Internet spacetime around them. The services they offer, and on which much of the Internet now depends, are integrated, centralized, and proprietary—the very opposite of what I tell my students.

Tejas N. Narechania and Scott Shenker's *How to Save the Internet* brings the stories we tell about the Internet back into line with the Internet as it actually is. Narechania is a law professor and Shenker a computer scientist. Their article is a seamless fusion of their expertise—and a cogent guide to the Internet's new normal and what it means for telecommunications policy and law.

How to Save the Internet begins with an overview of the traditional law-school description of the Internet, which is a model of clarity and economy. (I'm adding the article to the teacher's manual for my Internet-law casebook as a highly recommended primer.) The Internet knits together networks around the world by *layering* another virtual network on top of them; the Internet Protocol (or IP) standard that defines this global network relies on these smaller networks to transmit individual packets of data, so that any given message generally passes through several destination networks before reaching its final destination. Narechania and Shenker efficiently review both the technical fundamentals and the business terms on which Internet Service Providers (or ISPs) connect their networks. From this technical overview, they extract three core principles embodied by the traditional Internet design:

- *Neutrality* meant that individual ISPs could not effectively discriminate among the traffic passing through their networks: not based on its content, not based on the identity of its sender, and not based on the application or device sending it. "Network neutrality" is the decades-long attempt to turn this technical neutrality principle into a binding legal obligation.
- *Interconnection* meant that the Internet is a true network of networks, owned and operated by different entities, rather than a centralized monolithic, monopolistic network, as the U.S. telephone network under AT&T largely was. The IP standard provided the technical foundation for interconnection; economically, it was based on a system in which smaller ISPs typically paid larger ones to connect, while ISPS of approximately the same size carried traffic to and from each other for free.
- *Generality* meant that the Internet is open to any and all applications: email, file downloads, video calls, social media, video games, streaming audio, and more. Unlike an older network, which supported exactly those services that its provider offered (e.g., Comcast offers cable television, telephone, and home-security services as part of its non-Internet packages), the Internet is open to anything offered by anyone. The term "permissionless innovation" is sometimes used to describe this principle, but "generality" is better because it captures *why* such innovation matters.

This is where the traditional story stops. Plenty of legal debates can take place within this framework, including large swathes of the network-neutrality debates. But equally often, these features are taken for granted in legal circles. Scholars and students simply *assume* that a new application can be an idea today and a startup tomorrow—and showing up in police blotters and courtrooms the day after. That's how the Internet works, after all.

Except that, increasingly, it isn't. As Narechania and Shenker explain, while the Internet overall had a decentralized peer-to-peer design, many individual applications were built with a centralized client-server architecture. This approach can offer better performance and better security than a peer-to-peer design. If Spotify needs to serve more users, it can add more servers; if those servers are vulnerable to intrusions and attacks, it can hire more security engineers to secure them. All of this might seem like a Spotify problem, not an Internet problem–except that Spotify and other large application providers have increasingly been addressing their performance and security concerns in the network itself, rather than purely on their own systems.

The first big change is the rise of content delivery networks (or CDNs) that cache content on computers geographically closer to the users who need it. To take their example, the NBA *could* transmit game clips to users across the U.S. from its headquarters in New York. But that would mean the same video clip might need to travel from network to network across the country thousands of times. It would be much more efficient to send the clip once to a server in Los Angeles and send it to Los Angeles-area users from there (and so on for many other local regions). And thus, CDNs provide caching services: clusters of servers located near users, which are connected by the CDN's own private network.

Narechania and Shenker use the term "enhanced service provider" (or ESP, a nice play on ISP) to describe what CDNs have become. In addition to their caching services, they provide significant security benefits. (In particular, it is much harder to launch a denial-of-service attack on a CDN.) Some ESPs are vertically integrated; Google runs an immense private network to support YouTube and its other user-facing sites. Others are public-facing; Akamai provides large-scale services for customers who don't want to build their own secure CDNs.

ESPs, however, call into question the guiding design principles of the Internet:

- ESPs are emphatically not neutral. It's not just that they can pick their customers (often on the basis of willingness to pay, but sometimes on speech-related grounds). More fundamentally, their networks are designed to support different kinds of applications differently, with specific optimizations for streaming video, gaming, or other high-performance applications.
- ESPs are not generally interconnected with each other. They connect out to the public Internet to deliver content, to be sure, but each of them has its own entirely private network reserved for its own use.
- ESPs are not general. They offer discrete, integrated services. While they're built on top of a general-purpose resource—computational power—they don't sell it on an unbundled basis.

The result is an Internet that is increasingly concentrated among a few immense ESPs. The interconnected public portion of the Internet—what we call "the Internet" on the first day of class—carries comparatively less traffic, has less resiliency, and is more economically marginal. Narechania and Shenker fear that ESPs will impede innovation, both in developing new applications and in improving the Internet itself. It's a familiar tune of oligopoly and stagnation, played in a surprising new key.

What to do about it? Narechania and Shenker propose creating technical standards for ESPs' services, and particularly for an ambitious "InterEdge" design for modular networking services provided by the

ESP's server clusters. Regulators could then require ESPs to interconnect (along the lines of the interconnection mandate in the 1996 Telecommunications Act), require neutrality on the basis of content and speaker. The result, they argue, would be to catalyze a new generation of innovative applications, just as the Internet itself did back in the day.

I regret to say that I found *How to Save the Internet*'s recommendations for how to save the Internet less compelling than its diagnosis of the problem. The InterEdge is an appealing vision in some ways, but for now, it is a concept of a plan. The authors' prior technical work describing it is quite interesting, but much more of a slog for the reader who does not already have a firm command of networking architecture, and even there, the InterEdge remains somewhat abstract. But then again, that is what future work is for.

How to Save the Internet is compelling and highly informative. If you want to learn more on the technical side, computer scientists Pamela Zave and Jennifer Rexford cover similar issues in greater depth in their book The Real Internet Architecture: Past, Present, and Future Evolution. Narechania and Shenker's particular contribution is to show how these technical developments have significant legal and regulatory consequences. Their article is a must-read for Internet-law policy and Internet-law pedagogy.

Cite as: James Grimmelmann, *The Edge of Tomorrow*, JOTWELL (November 28, 2025) (reviewing Tejas N. Narechania & Scott Shenker, *How to Save the Internet*, __ **Berkeley Tech. L.J**. __ (forthcoming), available at SSRN (Mar. 18, 2025)), https://cyber.jotwell.com/the-edge-of-tomorrow/.