6-11-2012

Internet Policy Going Forward: Does One Size Still Fit All?

Christopher S. Yoo
University of Pennsylvania Law School, csyoo@law.upenn.edu

Internet Policy Going Forward: Does One Size Still Fit All?

Christopher S. Yoo

The Federal Communications Commission's (FCC's) December 2010 Open Internet Order represents a major turning point in U.S. regulation of the Internet, although the direction that Internet policy will take in the future is far from clear. The D.C. Circuit's 2010 decision holding that the FCC lacked the authority to sanction Comcast for interfering with BitTorrent and other peer-to-peer technologies1 ensured that the initial clashes over the Open Internet Order would focus on the scope of the FCC's jurisdiction.2 At the same time, public interest groups have challenged the FCC's decision to leave wireless broadband services outside the scope of some the rules.3 Resolving the merits of these particular controversies is an important matter that will likely represent the primary focus of Internet policy for the foreseeable future.

At the same time, the Open Internet Order is also based on a number of other propositions that may well prove even more influential over the long run and are worth exploring critically. One premise that I find particularly interesting is the belief that the Internet's past success stemmed from the fact that there has always been a single Internet that was open to everyone.4 As is always the case, this argument is only as persuasive as the assumptions on

---

[1] Comcast Corp .v FCC, 600 F.3d 642 (D.C. Cir. 2010).

[2] Verizon has filed an appeal with the D.C. Circuit challenging the FCC's jurisdiction to issue the *Open Internet* Order. *See* Verizon v. FCC, No. 11-1355 (D.C. Cir. filed Sept. 30, 2011).

[3] *See* Free Press v. FCC, No. 11-2123 (1st Cir. filed Sept. 28, 2011); People's Production House v. FCC, No. 11-3905 ag (2d Cir. filed Sept. 26, 2011); Media Mobilizing Project v. FCC, No. 11-3627 (3d Cir. filed Sept. 26, 2011); Mountain Area Information Network v. FCC, No. 11-2036 (4th Cir. filed Sept. 26, 2011); Access Humboldt v. FCC, No. 11-72849 (9th Cir. filed Sept. 26, 2011). On October 6, 2011, the Judicial Panel on Multidistrict Litigation consolidated all of these appeals in the D.C. Circuit. *In re* Petitions for Review of the Federal Communications Commission's *In the Matter of Preserving the Open Internet*, No. 1:11-ca-01356 (J.P.M.L. Oct. 6, 2011) (order granting motion to consolidate).

[4] Preserving the Open Internet, Report and Order, 25 FCC Rcd 17905, 17934 ¶ 49 (2010) [hereinafter Open Internet Order] ("[T]here is one Internet (although it is comprised of a multitude of different networks), and ... it should

which it is based. In this chapter, I would like to examine this proposition critically and explore the policy implications that might follow if it turned out not to be true. Exploring the ways that the Internet is likely to deviate from this "single Internet" vision and the implications of those deviations may provide greater insight into how best to shape Internet policy over the next five years.

The Longstanding Existence of Disconnected and Hybrid Networks

Despite the fact that discussions often somewhat casually treat the Internet as a single entity, the Internet is widely recognized by those familiar with its operation to be a “network of networks” composed of many smaller networks interconnected together. All of those networks exchange information through a uniform standard known as the “Internet Protocol” (IP), which represents the glue that binds the components into a larger coherent whole. In addition, Internet traffic has traditionally traversed backbone providers that exchanged information at public exchange points. When this was the case, networks exchanged traffic in a largely uniform manner.

Over time, an increasing amount of traffic has begun to deviate from the traditional pattern.5 Most importantly for our purposes, some firms rely on the Internet Protocol, but transmit their data over proprietary networks. For example, during the Internet’s early days, the acceptable use policy imposed on the NSFNET (the National Science Foundation Network)

---

remain open and interconnected....”); *id.* at 17956 ¶ 93 (“There is one Internet, which should remain open for consumers and innovators alike....”); *id.* at 18041 (statement of Chairman Julius Genachowski) (“There is one Internet, and it must remain an open platform....”).

[5] *See* Christopher S. Yoo, *Innovations in the Internet’s Architecture that Challenge the Status Quo*, 8 J. ON TELECOMM. & TECH. L. 79, 85-90 (2010).

prohibited uses unrelated to research and education. This prompted the creation of private backbones such as PSI, UUNET, and CerfNET, which operated in parallel with the NSFNET backbone without interconnecting with it.6 In later years, firms began providing IP-based voice services without traversing the public Internet in order to ensure improved security or greater guaranteed quality of service, which caused the FCC to refrain from referring to these services as "Internet telephony" and opt instead to use the terms "IP telephony" or "voice over Internet protocol" (VoIP).7 In addition, a large proportion of IP-based video passes via managed services without touching the public backbone.8 In short, a significant amount of IP-based traffic is travelling over various proprietary networks that bypass the public Internet altogether. A large amount of additional IP-based traffic employs hybrid systems that rely on proprietary or non-IP based technologies to form part of the connection.9 In addition, most of these networks engage in extensive network management that prioritizes certain traffic over other traffic.

Still other networks are moving away from exchanging data on the basis of IP in favor of other technologies that support greater functionality. The most popular of these is known as MultiProtocol Label Switching (MPLS). Instead of routing based on IP addresses, MPLS adds a label to the front of each packet and routes on the basis of that label. In addition, each flow (known as a Forwarding Equivalence Class) is assigned a specific path through the network. Information about the label and the associated route are propagated to other MPLS-enabled

---

[6] Michael Kende, *The Digital Handshake: Connecting Internet Backbones* 5 (Fed. Communications Comm'n Off. of Plans & Policy Working Paper No. 32, 2000), http://www.fcc.gov/Bureaus/OPP/working_papers/oppwp32.pdf; Barry M. Leiner et al., *The Past and Future History of the Internet*, COMM. ACM, Feb. 1997, at 102, 103.
[7] Federal-State Joint Board on Universal Service, Report to Congress, 13 FCC Rcd 11501, 11541 n.173 (1998).
[8] For example, AT&T's U-Verse network uses private networks to distribute video content to its central offices and then uses prioritized service to deliver programs to individual homes.
[9] For example, third-generation wireless devices employed hybrid systems that used legacy, non-IP based technologies to connect from the Internet gateway to the end-user device. Christopher S. Yoo, *The Changing Patterns of Internet Usage*, 63 FED. COMM. L.J. 67, 80, (2010). Many content delivery networks (CDNs) bypass the public backbone by using private connections to connect directly to their distributed caches. Yoo, *supra* note 5, at 88-90.

routers.10 Because labels are shorter than IP addresses, routers can direct traffic more rapidly. The fact that the route for a particular flow is defined in advance gives end users greater control over security. In addition, MPLS can support load balancing simply by dividing traffic between the same two endpoints into two separate Forwarding Equivalence Classes and assigning them different paths. Most importantly for the purposes of this chapter, in determining the particular path that a particular flow will travel, the MPLS router can match the quality of service demanded by the flow with the resources available along possible paths.11 MPLS can also recover more quickly from route failure. While once used exclusively within a single network, MPLS is now being employed across multiple networks, although doing so requires elaborate coordination between the networks to exchange the necessary information about the labels.

In short, a significant amount of data traffic relies on proprietary or hybrid networks rather than the public Internet. In fact, a leading industry study estimates that more than 25% of all IP-based traffic did not rely exclusively on the public Internet in 2010, and it forecasts those numbers to remain more or less stable through 2015.12 Almost all of these networks require the negotiation of private interconnection agreements and engage in extensive network management and prioritization. As such, they represent a significant deviation from the type of open, interoperable networks presumed by the FCC's Open Internet Order.

Moreover, it is plausible that positing the policy challenge as a mutually exclusive choice between managed and unmanaged networks may represent a false dichotomy. The routers within the core of the network are becoming increasingly programmable so that they can be

---

[10] Eric C. Rosen et al., *Multiprotocol Label Switching Architecture* (IETF Request for Comments 3031, rel. Jan. 2001), *available at* http://www.rfc-editor.org/rfc/pdfrfc/rfc3031.txt.pdf.
[11] William Stallings, *MPLS*, INTERNET PROTOCOL J., Sept. 2001, at 2, 3-4.
[12] CISCO SYS., INC., CISCO VISUAL NETWORKING INDEX 6 tbl. 2 (2011), *available at* http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf.

dynamically reconfigured from the default setting of providing traditional Internet service to provide a dedicated circuit instead. One leading example is Internet2's Interoperable On-demand Network (ION), which allows researchers to set up temporary or long-term dedicated point-to-point optical circuits to support large data transfers and other bandwidth-intensive applications.13 This represents a sea change from Internet2's previous position disfavoring network management as a solution.14

The ability to reconfigure routers dynamically illustrates the problem wigth framing the decision as an antipodean choice between two diametrical extremes. Instead, the two visions of a single, openly accessible network on the one hand and a completely Balkanized universe of noninterconnected private networks on the other represent polar extremes on a spectrum of possible policy responses. Only by understanding the countervailing considerations can one understand the forces that determine where along that spectrum social welfare would be optimized and what types of changes in the economic and technological environment might cause the optimal balance point to change.

Network Diversity vs. Network Size

The benefits attributed to the open Internet result in large part from the belief that being part of the largest possible network creates benefits for every part of the Internet community. End users benefit from being able to access content and applications without having to obtain

---

[13] Internet2, *Internet2 ION: User-Provisioned Dedicated Circuits Reserved in Advance or On Demand*, http://www.internet2.edu/ion/ (last visited Feb. 21, 2012).
[14] *Net Neutrality: Hearing Before the S. Comm. on Commerce, Science, and Transportation*, 109th Cong. 64, 66 (2006) (statement of Gary R. Bachula, Vice President, External Affairs, Internet2), *available at* http://commerce.senate.gov/public/?a=Files.Serve&File_id=c5bf9e54-b51f-4162-ab92-d8a6958a33f8.

approval from network providers. Content and applications providers can reach the entire universe of potential customers no matter how large or small they are.15 The increase in demand, in turn, encourages network providers to upgrade their infrastructure, driving what the FCC views as a "virtuous circle of innovation."16 This argument takes as given that adding more users to the network always creates additional benefits. Not only does it posit inexhaustible demand-side returns to scale; it also presumes that there are no countervailing considerations that could offset network size as a source of additional value.

Basic economic principles raise serious doubts as to the veracity of these propositions. As an initial matter, the basic principle of diminishing returns to scale dictates that the marginal benefits from adding additional customers will decline as the network continues to expand.

Moreover, the literature on network economic effects has long recognized that increases in network size may be less important if end users value certain connections more than others.17 As a general matter, most Internet users tend to focus their online activities on a relatively small number of endpoints. For example, studies have shown that although Facebook users have on average approximately 130 friends, they communicate directly with no more than four other people each week and only six other people each month.18 My own Internet usage exhibits similarly asymmetric usage patterns. When I connect through my PC at home, I interact with two network locations with greater frequency than any others: my email server and my office computer (via remote desktop access). Beyond that, I concentrate the bulk of my usage in a handful of websites that provide financial services (such as my bank) and news. While I sometimes access a broader range of information, such as when I plan a vacation, conduct

---

[15] Open Internet Order, *supra* note 4, at 17910 ¶13.
[16] *Id.* at 17910-11 ¶14, 17922 ¶28, 17927 ¶38, 17929 ¶40.
[17] *See* MILTON L. MUELLER, JR., UNIVERSAL SERVICE 20–24 (1997).
[18] PAUL ADAMS, GROUPED 23 (2012).

academic research, or search for activities for my children, I would willingly trade off somewhat less effective connectivity to the broader range of websites in order to obtain better connectivity to the locations that I know I visit the most. Although the ability to access other locations still has value to me, the fact that I visit them less often means that it has less value.

When the value that end users place on locations is heterogeneous, value depends as much on whether a particular network provides access to the locations that an end user values the most as it does on network size. The same is true for content and application providers. The fact that the revenue generated by the Internet overwhelmingly derives from advertising means that content and application providers will place a higher value on end users who possess desirable demographic characteristics and who can be targeted more effectively.19 Again, value is driven not by having access to more customers, but rather by having access to the customers who are the most valuable to a particular end user. Such differences in value represent a significant countervailing factor that can more than offset the benefits from raw network size.20

In addition to providing better connections to the most preferred locations, networks can also provide value by offering different services. Consider the way that the applications being carried by the network are changing.21 When the Internet first emerged, the overwhelming majority of traffic was dominated by two applications: email and web browsing. Although both applications are quite different in many ways, both are essentially file transfer applications for which application performance is determined by when the last packet arrives. As such, these applications are not particularly sensitive to variations in the arrival times of intermediate packets so long as they do not affect the arrival of the last packet.

---

[19] Christopher S. Yoo, *Architectural Censorship and the FCC*, 78 S. CAL. L. REV. 669, 683 (2005).
[20] Christopher S. Yoo, *Beyond Network Neutrality*, 17 HARV. J.L. & TECH. 1, 34-36 (2005).
[21] *See generally* CHRISTOPHER S. YOO, THE DYNAMIC INTERNET 19-36 (2012).

The applications that are now beginning to dominate the Internet are quite different. Specifically, Internet traffic is increasingly becoming dominated by video.22 In contrast to file-transfer applications such as email and web browsing, which are not particularly sensitive to irregularities in the spacing between intermediate packets, any variation in the arrival pattern of intermediate packets can seriously degrade video application performance. If the application can tolerate a brief delay before video playback begins, applications can minimize the impact of such variations by delaying playback until a sufficient number of packets are stored in a buffer to smooth out any variations in arrival rates. Buffering, however, is a poor solution for interactive applications, such as video conferencing.23

Although the Internet's protocols were designed to be as flexible as possible, no architecture can do everything well. Instead, every architectural design must pick and choose among the functions it will support. In so doing, it necessarily omits support for other functions required by other applications. Forcing everyone to interconnect in the same manner thus inevitably promotes applications optimized for the services supported by the particular network architecture while simultaneously inhibiting applications that require services that the network cannot provide.24 In particular, the current Internet provides service on a best-efforts basis that provides no guarantees about reliability or throughput times, which has long inhibited its ability to support real-time applications.25

The loss of variety associated with standardization thus represents one of the often ignored welfare losses associated with requiring that everyone be connected to a single, uniform

---

[22] *See* CISCO SYS., *supra* note 12, at 2, 9, 12-13.
[23] Yoo, *supra* note 9, at 70-72.
[24] Yoo, *supra* note 20, at 20-22; Yoo, *supra* note 9.
[25] Yoo, *supra* note 20, at 21-22; Yoo, *supra* note 9, at 85; Christopher S. Yoo, *Network Neutrality, Innovation, and Consumers*, 2008 U. CHI. LEGAL F. 179, 228-34.

network.26 In addition, diversifying service offerings also makes sense as a matter of business strategy. Business leaders have long recognized that one of the worst ways to compete in a market is to adopt a me-too strategy that simply mimics the actions of the market leader. Instead, leading books on business strategy exhort companies to try to differentiate what they are doing.27 Although such differentiation was once criticized as spurious and anticompetitive,28 we now recognize that product differentiation can create real benefits to consumers by allowing them to purchase goods that provide a better fit with their preferences.29 The welfare losses associated with the reduction in scale can be offset by the welfare benefits realized from product variety.

Simply put, the optimality of the decision to standardize on a single, universally available network depends on the heterogeneity of demand. If what everyone wants from the network is the same, the optimal strategy is to design a single, large network optimized in the manner that end users want. As the services that end users want from the network begin to diverge, network providers have strong incentives to diversify their offerings in response. Such responses are not necessarily anticompetitive and do not necessarily harm consumers. Instead, such diversification may be regarded as the natural byproduct of providers' attempt to satisfy consumer demand that is becoming increasingly varied. Indeed, such heterogeneity can make it possible for smaller providers to pursue niche strategies that allow them to survive despite their inability to fully realize the available economies of scale. Conversely, preventing firms from varying their

---

[26] *See* Joseph Farrell & Garth Saloner, *Standardization, Compatibility, and Innovation*, 16 RAND J. ECON. 70, 71 (1985) (noting that "reduction in variety" represents one of the "important social costs" of standardization); Michael L. Katz & Carl Shapiro, *Systems Competition and Network Effects*, 8 J. ECON. PERSP. 93, 110 (1994) (noting that "the primary cost of standardization is loss of variety: consumers have fewer differentiated products to pick from").
[27] CARL SHAPIRO & HAL R. VARIAN, INFORMATION RULES 26-27, 32-33 (1999).
[28] JOE S. BAIN, BARRIERS TO NEW COMPETITION 114-43, 204, 212, 216, 220, 223, 236, 239, 262, 279-81, 263-317 (1956).
[29] Yoo, *supra* note 20, at 30; *see also* Christopher S. Yoo, *Copyright and Product Differentiation*, 79 N.Y.U. L. REV. 212, 252-53 (2004).

services forces them to compete exclusively on price and network size, which are considerations that tend to favor the incumbents.30

This type of differentiation carries its own downsides. Those with strong preferences for the services offered by one particular provider would not regard the other options as substitutes, which would allow one provider to exert a degree of market power over them. Differentiation thus represents the type of tradeoff inherent in imperfect price discrimination, which effectively divides the larger market into submarkets and allows the provider to engage in monopoly pricing against each submarket. That said, it is important to keep in mind that the welfare implications are ambiguous, depending on whether the welfare benefits of expanding the number of customers served exceed the welfare losses associated with monopoly pricing in each submarket or vice versa.31 In addition, enabling some consumers to obtain services that fit better with their preferences creates welfare benefits from outside the price-quantity space of classical economics.32 The upside becomes more likely to dominate the downside as the ways that consumers use the network become increasingly varied over time.

Moreover, benefits from network size need not cause market failure if an adapter or converter exists that can allow end users on one network to interconnect on other networks. If the adapters function perfectly, end users will seem like they are on a single network, and the market will achieve the perfectly competitive result.33 Even if interconnection is imperfect, markets need not necessarily fail.34 Specifically, if the benefits of being part of a larger network are

---

[30] Yoo, *supra* note 20, at 34-35.
[31] Christopher S. Yoo, *Rethinking the Commitment to Free, Local Television*, 52 EMORY L.J. 1579, 1622-25 (2003).
[32] *See supra* note 29 and accompanying text.
[33] *See* Michael L. Katz & Carl Shapiro, *Network Externalities, Competition, and Compatibility*, 75 AM. ECON. REV. 424, 439 (1985) ("[I]f the costs of adapting are negligible, and there are no other entry barriers, the market will be perfectly competitive.").
[34] *See* Joseph Farrell & Garth Saloner, *Converters, Compatibility, and the Control of Interfaces*, 40 J. INDUS. ECON. 9, 32 (1992) ("[C]onverters might seem to promise the benefits of compatibility without the loss of variety due to

strong enough to induce standardization in the absence of converters but not strong enough to do so in the presence of adapters, the availability of imperfect adapters can reduce social welfare. Conversely, if the benefits of being part of a larger network are too weak to induce everyone to join a single network, imperfect adapters tend to increase compatibility and social welfare.35 The fact that prioritization and other forms of network management represent precisely this form of imperfect compatibility underscores that it is too facile simply to assume that the fact that networks increase in value as they increase in size necessarily leads to market failure.

Debunking the Myth of the One Screen

The need for a single network to which all industry actors have open access is also tacitly based on what I have called "the myth of the one screen": the assumption that end users will rely on a single Internet connection.36 The logic is simple. Only if end users maintain a single broadband subscription is it important that they be able to access the entire universe of content through any one connection. Conversely, if they maintain more than one connection, the fact that one of them offers better connectivity to certain content and applications is less problematic.

One of the growing changes in the industry structure is the growing tendency for end users to multi-home, that is, subscribe to more than one connection. Enterprise customers routinely contract with multiple providers to protect against loss of service. End users also routinely maintain both landlines and mobile phones to gain the advantage of mobility and to

---

standardization. Indeed, when converters are costless and perfect, this is true; but that is rare. In general, the welfare impact of the availability of converters is ambiguous.").

[35] *Id.*

[36] Yoo, *supra* note 21, at 122-26.

guard against network failure.37 I suspect that in equilibrium everyone will maintain at least two connections: one wired and one wireless. They may even optimize their usage to minimize cost and to take advantage of the technical differences.38

Multi-homing creates several advantages. The ability to obtain service from an alternative provider naturally restricts the market power that any one provider can exercise against any user.39 More importantly for our purposes, requirements that every connection represent all things to all people can be very costly. To take the most salient example, wireless broadband providers are currently struggling to support video services. Requiring that they support all applications on the same terms would require them to incur significantly higher capital expenses in order to offer service, thereby increasing the number of subscriptions that the provider will need to sell to break even. Given the severe restrictions on the available spectrum, this result will both increase prices paid by consumers and will make it impossible to provide service on an economical basis in marginal areas, thereby worsening the digital divide.

Private Networking as an Exit Option

The assumption that the Internet must remain open and universally interconnected also overlooks the fact that the imposition of government regulation rarely results in a stable equilibrium. Instead, the ultimate impact of such regulations can only be understood after the series of reactions and counter-reactions that the regulations are likely to stimulate are taken into account. Indeed, mandates of equal access to traditional telephone networks (along with attempts

---

[37] Yoo, *supra* note 5, at 86-88.
[38] *Id.* at 87-88.
[39] *See* Stanley M. Besen et al., *Advances in Routing Technologies and Internet Peering Agreements*, 91 AM. ECON. REV. (PAPERS & PROC.) 292, 292, 295 (2001).

to use rates for interconnection to implement other social policies) stimulated demand for private, bypass services that unless also regulated threatened to defeat the goals of the original regulation.40

Consider the likely impact of the fact that the Open Internet Order contains exceptions for purely private networks, networks that provide limited functionality, and specialized services.41 Providers subject to the Open Internet Order thus have strong incentives to bypass the public Internet – and the accompanying regulations – simply by shifting their resources to private networks. The problem is that shifting to dedicated networking resources eliminates one of the central efficiencies created by the Internet by preventing multiple users to share resources. In particular, network operators who need quality of service assurances will generally find it most efficient to offer their bandwidth to others only on a secondary basis, available only when the primary owner does not need it. Indeed, this is the central motivating principle underlying the public safety network envisioned for the D-block in the 700 MHz auction42 as well commercial offerings such as AT&T's U-Verse and proposals for more efficient use of spectrum.43

Thus, those who advocate mandating open access in the name of preserving a unified Internet must consider the possibility that such a mandate might actually create incentives toward greater fragmentation. In addition, restricting network owners' ability to share surplus capacity threatens to increase the cost of broadband where it is available, while at the same time worsening the digital divide by reducing the geographic areas in which such service can break even.

---

[40] *See* Dennis L. Weisman, *The Proliferation of Private Networks and Its Implications for Regulatory Reform*, 41 FED. COMM. L.J. 331 (1989).
[41] Open Internet Order, *supra* note 4, at 17928 ¶39, 17933 ¶47, 17951 ¶79, 17965–66 ¶¶112-114.
[42] *Service Rules for the 698–746, 747–762 and 777–792 MHz Bands*, Second Report and Order, 22 FCC Rcd 15289, 15431 ¶396, 15437 ¶414, 15437 ¶416 (2007).
[43] *See* Jon M. Peha & Sookan Panichpapiboon, *Real-Time Secondary Markets for Spectrum*, 28 TELECOMM. POL'Y 603, 604 (2004).

Conclusion

The foregoing analysis raises two separate critiques to the claims that openness is a critical aspect of the Internet's past success that must be preserved. Any assessment of such a categorical claim must grapple with the extent to which the Internet was, in fact, unified in the past and with whether the circumstances that influenced the Internet in the past remain relevant to the Internet's future. Moreover, the unqualified manner in which this "single Internet" claim is advanced obscures the reality that every important policy decision typically involves a tradeoff between competing considerations. Most relevant for our purposes is the fact that mandating open interconnection implicitly presumes that the exclusive source of value to end users is raw increases in network size. Framing the issue in this manner fails to consider that end users typically place a premium on being able to reach a small number of locations and run a discrete number of applications. The policy analysis must also take into account that end users are likely to maintain more than one connection and that any attempt to mandate open access is likely to provoke a series of reactions and counter-reactions that may frustrate the goals of the initial regulation.

On a broader level, policymakers would benefit from taking seriously the possibility that the days of a "one size fits all" approach to Internet regulation may well be over and that looking backwards for the lessons of the past may not always be the best way to promote future success. As anyone can attest who experienced how quickly AOL keywords shifted from the critical way to access customers to near-complete irrelevance, the technological and economic environment surrounding the Internet is constantly undergoing rapid, dynamic change. The policies developed

for a world dominated by PCs using cable modem or DSL service in which the browser was the critical platform may no longer be the right framework for a world increasingly dominated by smart phones attached to wireless broadband networks in which the critical platform is now the app store and the wireless operating system. The growing heterogeneity of the technologies, end user demands, and business relationships underlying what is now often referred to as the Internet ecosystem may require reframing the issues in a fundamentally different manner.