Policy Brief

Stanford Institute for Economic Policy Research

Will Public Policy Protect the Architecture of the Internet?

Paul A. David

Everyday life in the world's economically advanced regions has been touched and in some parts substantially transformed by the advent of the Internet. The expansion of scale that this system has achieved in so brief a time is breathtaking. The Internet can be regarded as the largest artifact in the known universe: There are now over 100 million network hosts, some 200 million PCs connected online, and almost 30 million Web sites on the World Wide Web. The pace of growth in global connectivity and the phenomenal proliferation of diverse innovations in applications software are marvels that distinguish this communications infrastructure's performance from that of its historical predecessors such as the telegraph and telephone networks.

The user perceives the Internet as though it were a single homogeneous network, but in actuality it is a softly integrated heterogeneous "network of networks." The openness and transparency of this "connection-less" communications system are properties derived from the distinctive "end-to-end" design of its architecture and transmission control mechanisms. These features enable the Internet to tolerate extreme diversity in the technical specifications of its constituent networks and platforms. That has made joining the system cheap, and highly attractive to new network operators, Internet Service Providers (ISP) and users. The transparency of the Internet's "end-to-end" architecture affords a particularly accommodating platform for developers of applications innovations. Software can be designed to run on the computers situated at the network's "edges" – taking data inputs and generating data outputs that traverse the intervening channels – without having to pay attention to the specifics of the computer hardware and software that perform the routing functions of the communication system.

March 2002

The technology of this infrastructure is not static. Fundamental alterations to support extensive peer-to-peer computing are, perhaps, more than a decade away; but strong pressures are mounting for more immediate engineering changes in the core of the network. The variety of "adaptive network modifications" that presently are under active consideration should not, however, be construed as obvious steps in some automatic process of technological optimization leading to an enhanced version of the Internet we know and love. Unless they are subject to independent expert assessments carried out within explicit public policy guidelines, some among the proposed engineering modifications would alter key performance features of this communications system. Whatever gains in social and economic welfare are to be expected from these "improvements," the full extent of their implications should be understood, so that promised direct gains in performance may be properly weighed against the possible losses to stakeholders of other social and economic benefits – particularly those deriving from the transparency of the Internet's inherited "end-to-end" design of the network's architecture.

At present, however, there is nothing in the political economy of the Internet to assure that assessments of this kind will be carried out, or that they will influence the engineering steps that are taken in reaction to perceived drawbacks of the present architecture. Many of these deficiencies are not new. They appeared quickly after the Internet was thrown open to general public and commercial traffic in the mid-1990s. The most salient among them are the difficulties of blocking delivery of unwanted content ("spam" or offensive material), suppressing malicious actions (e.g., "viruses") and pricing bandwidth to reduce transmission delays.

Technical remedies for some of these already are being implemented by the introduction of so-called filters installed in "firewalls" at the edges of the network. But the latter also are being deployed by third parties that can act without the users' consent: According to a recent report, the government of China has been able in effect to "firewall" the entire country, thereby controlling connections with the rest of the Internet in addition to monitoring the content of internally generated traffic. What makes this feasible for an authoritarian government and a business corporation alike is that there are a relatively small number of paths connecting its domain to the rest of the network; the same would be true for an ISP. Inserting firewalls and filters at those few passage points is an effective and comparatively low-cost means of imposing selective controls on the messages that residents of the domain are able to exchange with the rest of the world. Equally, it allows the insertion of clandestine traffic analysis and content monitoring by outside parties. This possibility gives rise to understandable concerns that – especially in the post-September 11 climate – ISPs may find it difficult to resist requests from government agencies to permit this to be done in the interests of "security."

The insertion of technical devices to enable governments to exercise control functions for political purposes, or to protect the integrity of the communications system, is only one way in which the original architectural features of the Internet may be compromised. There are economic incentives for ISPs to adopt engineering innovations that would support high-value data transport services – services for which the precursor networks forming the Internet were not designed. The Internet's TCP/IP protocols – which provide capabilities for reassembling the data packets in proper order, re-transmitting lost packets and confirming complete delivery – offer a "best effort" quality of service. While this has been successful in supporting a wide range of applications, it does not establish a dedicated connection between the sender and the receiver; and so, it cannot make any guarantees for users as to when, or even whether, a message will be delivered. Network services like email and Web browsing easily tolerate the existing transmission delays and delay variations that are characteristic of the TCP mechanism, but these fatally degrade voice telephony and video services over the existing Internet.

Consequently, would-be vendors of voice telephony and real-time video on the Internet and other complementary services have a keen interest in proposals to modify the layer of technology that controls and manages flows of data-packets, in order to achieve a "quality of service" approximating that of the public switched telephone network. That would entail modifying the Internet's routers in ways that terminal hard-

ware and software would need to recognize and take into account; it is perhaps the most likely of the plausible evolutionary paths along which the ending of end-to-end architecture would be driven by private business initiatives.

Another source of pressure upon the Internet's architecture comes from the enlarged scope for business strategies premised upon exploiting opportunities for "regulatory bypass," which in itself has posed new challenges to the ability of public authorities to effectively regulate telecommunications industries. In the United States, network operators in the long-regulated telephone business that offer broadband access to the Internet have been required, largely for reasons of competition policy goals, to provide their customers with open and nondiscriminatory access to other broadband ISPs. Cable companies, on the other hand, although performing exactly the same functions, find themselves under weaker regulatory constraints in this regard. This leaves the way open for *some* ISPs to pursue a strategy of creating what might be described as "restricted access shopping precincts in cyberspace": islands on the Internet where subscribing customers will be offered particular, pre-selected bundles of communications services, information applications, auctions and other e-shopping opportunities, as well as games and databases.

Such a strategy obviously could create sources of indirect profit for the "cyber-mall landlord": In the absence of multiple avenues for Internet access, the firm in question could exercise considerable market power vis-à-vis the originators of the variety of goods and services being offered there. Even though restricted cyber-malls of this kind might well prove attractive for many customers seeking convenient and low-cost access to standardized packages of regularly upgraded information goods and services, the adverse long-run effects upon competition would reduce the economic benefits they derived. Significantly, in the scenario just envisaged, the effectiveness of the Internet as a platform for innovation would be curtailed, restricting Net-wide innovations. Thus, it is not surprising that some informed observers have expressed dismay that the existing regime of regulation (and non-regulation) in the United States may permit the cable companies to bundle broadband access with selected application service offerings.

Rational discussion of the tradeoffs among diverse policy options – preserving a transparent platform for technical and business innovation, or making the network unavailable for use in coordinating terrorist attacks, or providing the Internet's users with voice telephony and "freedom from spam" – requires an elevated measure of public awareness of the system-level implications of specific proposals for regulatory or engineering modifications.

Without such understanding, it would be particularly difficult to protect the special performance features of the Internet as a public domain for "information discovery," for the creation of new modes of discourse and democratic participation, and for the facilitation and coordination of learning, research and innovative activities that respond to the needs and desires of diverse communities distributed throughout the world.

Further Reading

"The Beginnings and Prospective Ending of 'End-to-End': An Evolutionary Perspective on the Internet's Architecture" (http://siepr.stanford.edu/papers/pdf/01-04.html)

The Stanford Institute for Economic Policy Research (SIEPR) conducts research on important economic policy issues facing the United States and other countries. SIEPR's goal is to inform policy makers and to influence their decisions with long-term policy solutions.

With this goal in mind SIEPR policy briefs are meant to inform and summarize important research by SIEPR faculty. Selecting a different economic topic each month, SIEPR will bring you up-to-date information and analysis on the issues involved.

SIEPR Policy Briefs reflect the views of the author. SIEPR is a non-partisan institute and does not take a stand on any issue.

For additional copies, please see SIEPR website at: http://SIEPR.stanford.edu

About the author

Paul A. David is a Professor of Economics and Senior Fellow at the Stanford Institute for Economic Policy Research, as well as Senior Research Fellow of All Souls College, and Professor of Economics and Economic History at the University of Oxford. An elected Fellow of the International Econometric Society, the American Academy of Arts and Sciences, and the British Academy, David is known internationally for his analyses of the role of "path dependence" in economic processes, which feature prominently in his contributions to American economic history, economic and historical demography, and the economics of science and technology. A strong focus upon contemporary public policy issues characterizes David's recent publications. David has published more than 120 journal articles and contributions to edited volumes. He is the author of Technical Choice, Innovation and Economic Growth, and other books, with several new works scheduled to appear during 2002-2003.

SIEPR **Policy Brief**

A publication of the
Stanford Institute for Economic Policy Research
Stanford University
579 Serra Mall at Galvez Street
Stanford, CA 94305
MC 6015