

The InterNAT: Policy Implications of the Internet Architecture Debate*

Hans Kruse¹

J. Warren McClure School of Communication Systems Management
Ohio University

William Yurcik²

Department of Applied Computer Science
Illinois State University

Lawrence Lessig³

Professor of Law
Stanford Law School
Stanford University

Abstract

In 1981, Saltzer, Reed, and Clark identified "end-to-end" principles related to the design of modern layered protocols. The Internet started out as a network in which all "intelligence" was placed in the end-nodes (hosts), while the network is strictly concerned with the best-effort delivery of individual packets. To an application residing on several hosts the network is therefore "transparent" in that it has no effect on the application other than facilitating the delivery of information between the applications. The Internet today is not as transparent as Saltzer et al. had envisioned. While most of the intelligence remains concentrated in end systems, users and network operators are now deploying more sophisticated processing within the network for a variety of reasons including security, network management, E-commerce, and survivability. For example end-users are deploying Network Address Translators (NATs) to circumvent problems related to IP address allocation, and firewalls and proxy servers for security at the interface between the user's network and the Internet. Network operators use packet filters and application level gateways to deal with security issues ranging from "spam" to denial of service attacks. In addition, network operators are deploying router software to enable differentiated levels of service, and to create virtual overlay networks for corporate clients. Each of these implementations removes a certain amount of transparency from the network by introducing "layer violations", i.e. access to non-network layer information inside the network. Applications and application-layer protocols have been found to react in unexpected ways to the presence of these layer violations. We note that a transition to IPv6 is a possible solution to the address allocation

* supported in part by a grant from NASA #NGT-30019 and the John Deere & Company

¹ author for correspondence; contact information: hkruse1@ohiou.edu telephone/fax (740) 593-4891/4889
hardcopy mail J. Warren McClure School of Communication Systems Management, 9 South College Street,
Room 197, Athens, OH 45701-2979 USA

² Telecommunications Program within Department of Applied Computer Science, Email: wjyurci@ilstu.edu

³ Email: lessig@pobox.com

issue, and it may slow down the proliferation of NATs; however, it is quite clear that layer-violating devices will be a permanent part of the Internet.

A true end-to-end model makes the Internet transparent and thus a commodity; in this scenario network operators compete based on price, bandwidth, and reliability. Outside the known issues related to facilities based carriers, there is little opportunity for anti-competitive behavior. The deployment of layer-violating network devices is straining the end-to-end model and creating a different competitive landscape. Given the large installed base of layer-violating network devices already within the Internet and recent denial-of-service attacks accelerating demand, Internet Service Providers have had to control traffic and protocols out of technical necessity. In a truly transparent network, the network operator is unaware of the applications being run by the connected hosts (security purists would argue that this is the desirable state in any communications network). In the presence of layer-violating devices, the network operator has to take explicit steps to enable end-user applications, usually by deploying gateways that mitigate the impact of the layer violations. This creates a clear opportunity for the network operator to engage in the enabling or disabling of applications on the basis of non-technical decisions, including the ability for the operator to engage in anti-competitive behavior.

In this paper we describe specific examples of the technical and policy problems caused by the introduction of this new processing within the network which is counter to the end-to-end Internet model proposed by Saltzer et. al. We describe a number of possible scenarios for anti-competitive strategies and argue that technical decisions that shape the Internet architecture may indeed render it more subject to legal and regulatory control. We conclude that the presumption should be in favor of preserving the architectural features that have produced the extraordinary innovation of the Internet while warning that a market failure may be occurring under the guise of technical pretenses.