

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

* 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

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 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.

1.0 Introduction

There are two classic models for intelligence within networks. [LEAR00] The first is an end-system model. Under this design, end-devices have no intelligence; the network devices to which they connect provide all the services. The telephone system is an example of just such a network. End-systems have well known benefits and costs. The absence of intelligence in end-devices makes them inexpensive to manufacture and manage; but network devices (central-office switches) in turn become expensive and complex to maintain.

The second model for intelligence within networks is end-to-end.⁴ First proposed by Saltzer, Reed, and Clark in 1981, the end-to-end model is a set of architectural principles that guide the placement of functions within a distributed system. [SALTZER81] According to the argument, lower layers of a distributed system should avoid providing functions that can be implemented in higher layers (end-systems), especially if (1) the function cannot be completely implemented in lower layers and (2) all applications would not benefit from such functions.

End-to-end thus shifts intelligence in a network to the application hosts. It therefore also shifts cost and management complexity from routers and switches to hosts. This is a benefit for those maintaining the network. Another benefit is that congestion control can be managed between hosts; it is therefore not required that state information be kept within routers to optimize performance.⁵ In the end-to-end design the network simply acts as a transparent transport mechanism for individual packets with each packet being labeled by a globally unique source/destination addresses.

⁴ We use the term “end-to-end model” while acknowledging that the original authors prefer to use the terminology “end-to-end arguments”

This notion of “transparency” implicit within an end-to-end design has a number of technical and policy consequences. It demands that network devices between two end-systems not modify information within the packet above the data link layer, except under well-defined circumstances.⁶ Changing IP addresses is not viewed as acceptable, nor is any change to the Transport layer (layer 4) or above.

2.0 The Problem: Unexpected Protocol Interactions

“The New York City Board of Education is using network address translators as a security measure to keep their 1000+ schools off the public network (Internet). Teachers are reporting that the networks are unusable because of them. Many of the educational benefits that the schools want to gain from being connected to the Internet are inaccessible because of the limitations network address translators place on the type of connections that may be made (and accepted).”⁷

The end-to-end model is a choice, not a necessity in the Internet’s design. The Internet Engineering Task Force (IETF) has traditionally been instrumental in supporting end-to-end with “rough consensus and working code.”⁸ In fact, one of the authors of the original end-to-end model paper, David Clark, chaired the Internet Activities Board (IAB) overseeing the IETF from 1981 to 1989. But the design is increasingly under threat. In reflecting on the state of the Internet in late 1999, a current member of the IAB

⁵ Noted Internet researcher Van Jacobson is quoted as stating, “Very simple. A router has only three choices when presented with a packet. It can transmit the packet. It can delay (queue) the packet. Or it can throw the packet away.” [CHEN98]

⁶ Network devices routinely update a "hop" counter in the network layer, and may record the route taken by a packet. Routers may also alter the content of one or more quality of service label fields in the network layer portion of the packet.

⁷ Jeffery Altman, Email posted to the IETF mailing list, December 1999.

⁸ motto of the IETF

and present/past chair of numerous IETF working groups, Steve Deering,⁹ summarized his thoughts on intelligence within networks with a slide - “Internet is Losing?”¹⁰ The examples he used included:

- unique IP addresses are no longer necessary
- the Internet is not always on (many users log-on via American On-Line etc.)
- end-to-end transparency is often blocked behind network address translators and firewalls

While the intelligence in the existing Internet remains concentrated in end-systems, users are increasingly deploying more sophisticated processing within the network for a variety of reasons including security, network management, E-commerce, and survivability. The following are some specific examples:

- The use of network address translators to solve IP address depletion problems.
- The use of performance enhancing proxies to tune protocols on links with unusual characteristics, e.g. in terrestrial and satellite-based wireless systems.
- The use of tunneling and other virtual private network techniques to provide secure connectivity over the Internet to an organization’s intranet/extranet.
- The use of firewalls and intrusion detection to prevent and respond to malicious attacks.
- The deployment of quality-of-service mechanisms to provide delay, delay jitter, and packet loss guarantees to applications and network services.

Our point is not to question the need that has led to each of these devices. No doubt each addresses important network needs that demand resolution. Rather than debate the benefit of each such device and their legitimacy with the network, we accept the notion

⁹ Steve Deering is also the inventor of IP multicast and lead designer of the next generation Internet Protocol (IPv6).

¹⁰ Closing Talk of *Networked Group Communications Conference* (NGC’99), Pisa Italy, Nov. 19, 1999.

that such technologies are here to stay at least for the short-term. End-to-end will not be reestablished by banishing such devices. But the presence of such devices threatens the existing end-to-end Internet. The problem is exacerbated in the present context by the inability of hosts and applications to detect intelligent network devices.

While the next generation of the Internet Protocol, IPv6, has been designed to solve many of these problems, migration will take time. Not only must protocol stacks and routers be upgraded, but applications must be changed to conform to the new structure of the IPv6 addresses. The good news is that IPv6 has been designed so that IPv4 and IPv6 can coexist while IPv6 is deployed gradually. The bad news is that this coexistence means that there will be less pressure to move from IPv4 to IPv6. In addition, devices designed to provide network security, improved performance, and network monitoring will continue to "break" the end-to-end model even in the IPv6 network.

In the meantime, there are significant costs to the end-to-end Internet. Applications and application-layer protocols interact in unexpected ways with intelligent network devices within the current IPv4 Internet model. This is a consequence of intelligent network devices reducing the transparency of the network. The critical element of transparency is some ability to predict how the network will behave.[CHEN98] To quote from a 1998 paper from the original authors of the end-to-end model:

*“Since lower-level network resources are shared among many different users with different applications, the complexity of potential interactions among independent users rises with the complexity of the behaviors that the users or applications can request. For example, when the lower layer offers a simple store-and-forward packet transport service, interactions take the form of end-to-end delay that can be modeled by relatively straightforward queuing models. **Adding priority mechanisms (to limit the impact of congestion) that are fixed at design time adds modest complexity to models that predict the behavior of the system. But relatively simple programming capabilities, such as allowing packets to***

change priority dynamically within the network, may create behaviors that are intractable to model....”[CHEN98]

The network therefore faces an important trade-off. Maintaining the largest degree of network transparency constrains interactions among different users of a shared lower level so that network behavior can be predicted, but it also creates performance and security problems that cannot be easily solved. Therefore, deployment of some of the features required by the network users requires the use of devices that violate end-to-end transparency; this diminishing transparency increases unexpected interactions between protocols.

We have identified three distinct types of unexpected protocol interactions that have been introduced by the diminishing of transparency due to the deployment of intelligent network devices:

- Some network devices attempt to read or modify portions of transmitted packets which the sending system assumes fixed. [i.e., performance enhancing proxies, network address translators]
- The use of IP tunnels creates the design issue of how to construct the second “outer” IP header upon tunnel ingress¹¹, and the more complicated issue of whether the original “inner” IP header needs to be modified upon tunnel egress,¹² based on changes that intermediate nodes made to the outer header.[i.e., tunneling and virtual private networks]
- Some devices on purpose, or due to limits in their design, prevent some packets from traversing that device. [i.e., firewalls, intrusion detection]

In this paper we examine these protocol interactions in a effort to understand recent protocol design decisions and their effect on the transparency provided by the end-to-end

¹¹ **ingress** is a path going **into** a network

¹² **egress** is a path **exiting** from a network

Internet model. We are particularly interested in examining the protocol structures involved to determine why the traditional protection against protocol interactions inherent in the layered protocols could not prevent the observed problems. The remainder of this paper is organized as follows: Section 3 describes the use and interaction of network address translators and other network devices which destroy transparency. Sections 4 states the policy implications of these challenges to the end-to-end Internet model. In Section 5 we close with a summary and directions for future work.

3.0 Layer-Violation Network Devices

One of several design philosophies behind the Internet protocols is to provide a variety of services based on the concept of layers. [CLARK95] This layered design is intended to provide needed information to each type of network device independent of the information required for other devices. A network device should normally operate at or below the network layer of the protocol stack, e.g., the IP layer in TCP/IP. End-systems rely on the end-to-end Internet model to provide transparency to layer information (IP layer and above) such that this layer information remains unchanged or invisible. However, a number of special circumstances have led to the creation of layer-violation (LV) network devices that rely on information from protocol layers they would not normally access. [KRUSE99]

3.1 Network Address Translators (NATs) / Performance Enhancing Proxies (PEPs)

“There is no longer a single Internet address space. We’re going to have to call it the InterNAT”¹³

¹³ Lloyd Wood; IETF mailing list February 17, 2000; Wood goes on to say, “Hey, that means the Internet is InterNATIONAL.”

NATs allow the use of private IP addresses in a private intranet while maintaining connectivity to the Internet through one or more global IP addresses. Since many applications assume that the end-system address is globally unique, NAT usually require application level gateways which modify application-specific sections of the packet where the end-system address has been embedded. These gateways cause changes in the packet that are unanticipated by the end-systems.[HAIN00, HOLDREGE00] A Network Address and Port Translator (NAPT) cannot forward a connection request from the Internet to a private network unless an administrative mapping has been provided for the port requested in the incoming packet. Other packets may be dropped or misrouted because the NAPT does not have the appropriate application-level gateway and thus fails to make corrections in the packet to allow the application's peer to respond.

It should also be noted that with the advent of dial-up Internet users whose IP address is allocated at dial-up time, the actual IP addresses of such users is purely transient. During their period of validity they can be relied upon end-to-end but these IP numbers have no permanent associations with the domain name of any host and are recycled for reuse at the end of every session. Similarly, LAN-based users typically use DHCP¹⁴ to acquire a new address at system restart.

PEPs are used in networks with unusual link characteristics.[ALLMAN99] These proxies may attempt to read transport-level information in the packet or they may add and delete packets from the flow. Many of these proxies can be bypassed by flows that do not permit such interactions, at risk of suffering from poor performance. Both NAT and PEP devices vastly complicate the deployment of IP-level security between end-systems [KRUSE99], and they may cause other failures that can be difficult to diagnose [CARPENTER00]. For instance, both the NAT and PEP devices usually do not report the fact that they either failed to correctly handle a packet, were bypassed, or dropped a packet they could not process due to insufficient information. Encrypted packets will be examined by the security software at the receiving end where modifications made by the

¹⁴ DHCP is the Dynamic Host Control Protocol

NAT or PEP device will be interpreted as illegal tampering and the packet will be discarded by the security software. While dropping packets is an auditable event, the sender of the packet is usually not notified.

3.2 Tunneling and Virtual Private Networks (VPNs)

IP tunnels are defined as a section of the network in which IP packets are encapsulated inside a second IP header (often called the “outer header”). The tunnel is designed to transport packets between two intermediate points in the network, without making reference to the actual IP packets during the tunnel section of the packet’s path. Tunnels can serve a number of purposes including:

- Transport of multiple protocols over an IPv4 router infrastructure (i.e., IPv6, IPX, AppleTalk) as well as service types not supported by intermediate nodes (i.e., multicast backbone or MBONE).
- Tunnels provide secure passage between two nodes at the edges of trusted domains. Inside the tunnel, original IP packets are encrypted and therefore completely inaccessible.
- Creation of VPNs. In this scheme, packets between two sites are carried over IP tunnels to provide isolation from the addressing and routing requirements of the Internet. A similar type of tunnel can be used to connect an off-site user to the corporate network.

The use of tunnels creates specific types of protocol interaction problems. Specifically how should the outer IP header be constructed at the tunnel ingress point? In general it seems reasonable to copy fields from the original IP header, however, this is not always the correct approach. In networks that provide quality of service control through resource reservation [TERZIS00] or differentiated service [BLAKE98], the tunnel may be used to traverse a portion of the network that cannot provide these services, and therefore requires that some of the original IP settings not be copied. In

other cases [FLOYD00], the ability of the tunnel egress point to provide certain types of processing will determine how to construct the outer IP header.

By far the more complicated issues arise upon tunnel egress. Some portions of the outer IP header may have been modified during tunnel traversal. Examples include updating of header fields that mark the packet as being in a particular differentiated service group, or updating of fields designed to provide explicit congestion notification to end-points. The tunnel egress node must merge the original IP header with the – possibly modified – outer IP header. The rules for doing this are ambiguous and different procedures may emerge. For example, from a performance and application perspective, one would wish to propagate congestion notification information across security tunnels. From a security perspective, one may wish to discard the entire outer IP header, regardless of its content, to prevent attacks based on the ability of hostile systems to modify the unprotected outer IP header inside the tunnel.

3.3 Firewalls, Intrusion Detection, and IPsec

Several devices can discard packets before they reach the end-system destination address. Most prominently, firewalls are designed to do just that for all packets that have not been entered in a permission list. Firewalls, by their very nature, fundamentally diminish transparency. Typically the source is not notified of the fact that the packet was dropped (although auditing of dropped packets can be performed at the firewall). In order to prevent attacks, many corporate firewalls will not permit network management packets (i.e., ICMP) to pass through.

Intrusion detection (ID) is a monitoring and auditing system for attempted and successful system breaches with the goal of detecting and ultimately preventing such activity. Because many attacks can be recognized by their signature (headers), the best place to process information is at the network layer. Since ID is based on algorithms which correlate network layer information with signatures, they require large amounts of storage. The state-of-the-art is reactive off-line processing. ID systems are currently

maturing and the next generation will rely on integration with routers to proactively monitor activity in real-time. One example of a transparency issue related to ID systems is fragmentation. While fragmentation is a useful method for supporting various media on internetworks, it may mean caching packets at the ID system to reassemble for inspection – a process that destroys transparency and could be a performance bottleneck.

At the other end of the spectrum from filtering and correlating packets is security. IPsec is actually an architecture - a collection of protocols, authentication, and encryption mechanisms – as described in [KENT98]. The loss of transparency is both a bug and a feature from the security standpoint.[CARPENTER00] To the extent it prevents the end-to-end deployment of IPsec, it damages security and creates vulnerabilities. For example, if a NAT is in the path, the best that can be done is to decrypt and re-encrypt IP traffic in the NAT with the traffic momentarily in plaintext. Noting NATs are prime targets for attack already, this is unacceptable. Indeed, NATs break other security mechanisms as well, such as Kerberos and DNSSEC, since these rely upon address values. In a weaker sense, the loss of transparency at an Intranet/Internet boundary may be considered a security feature since it is a well-defined point to enforce security policy. However, such a security strategy is vulnerable to insider attack and boundary penetrations which expose the entire intranet to trivial attack. Lastly, where cryptographic algorithms are used, protocols should be designed to permit alternative algorithms to be used. There have been several efforts by corporations to embed their own patented cryptographic algorithms within a protocol to capture a market while at the same time severely limiting end-to-end transparency.

Electronic commerce applications commonly require the implementation of procedures to insure confidentiality (usually via encryption), authentication, and non-repudiation, and availability (especially the prevention of denial-of-service attacks). In many cases these systems are based on end-to-end semantics that rely on the transport layer information remaining unchanged within the network. LV network devices also hinder the deployment of this security infrastructure.[KRUSE99]

3.4 Quality-of-Service (QoS) Mechanisms

Classically, the end-to-end model views the network as a monolithic entity that provides a single QoS to all users, best-effort delivery. The Internet has expanded to incorporate applications with requirements for guarantees on network behavior beyond the best-effort delivery. In the case of mechanisms such as RSVP, the host signals to the network the level of service it requires, whereas with differentiated services (DIFFSERV) the network prioritizes traffic without the host's knowledge or consent.[BRADEN97, BLAKE98] Both end-systems and LV network devices cooperate to provide deterministic and statistical QoS guarantees on metrics such as delay, delay variation, and packet loss rates. It is an open problem how to provide QoS guarantees over the Internet but the proposed schemes incorporate LV network devices which will introduce unexpected protocol interactions.

4.0 Policy Implications

Technical architectures embed policy choices. Changes in those architectures change those policy choices. The architecture of end-to-end is no exception. The issues that end-to-end affect reach beyond the value of any particular LV technology. Compromising end-to-end thus creates externalities. Changes may render the Internet more subject to private or state control. [LEMLEY99] Successful analysis of such changes therefore demands an ability to synthesize both the technology and policy issues.

Our aim in the sections that follow is to identify some of the policy issues implicated by changes in the end-to-end architectures. We begin by identifying analogies to end-to-end architectures in other contexts of social policies. We then consider the social risks that changes might present.

4.1 End-to-End Analogies

There is an analogy between the values that an end-to-end architecture embeds, and the values exemplified in other familiar social systems. The essence of this value vests power in end users. End-to-end is a structure for assuring the bottom-up control over the evolution of the Internet. Like a competitive market setting prices, or a federated republic preserving free trade, end-to-end minimizes both the control a central actor might have, or the opportunity for control that particular individual actors might have.

These two effects are distinct, and their difference can be seen in relation to the constitutional design of the American republic. Like an end-to-end system, the U.S. constitution minimized the power vested in a central authority. (The states, and the People, retained all power not granted to the federal government; the power originally granted was minimal.) But not all choices were originally left to the states. In some respects, states under the original design were constrained. In particular, states were restricted in their power over interstate commerce. A state could not, consistent with the constitution, discriminate against commerce flowing from outside its borders; it could not burden interstate commerce more than necessary to achieve legitimate state ends.

These restrictions on state freedom serve an important national goal. They assure the free flow of commerce within the United States, and thereby spur greater and more diverse commerce. Like the effect of end-to-end, this guarantee of a neutral market induces more innovation; innovators know that their efforts will have the benefit of a large national market. Both the commitment to a decentralized regulatory regime, and a restriction on the scope of that regulatory regime therefore advance the national market.

End-to-end in network design functions in a similar way. By pushing “intelligence” in the network to the application layer, the system decentralizes control over network use and functionality. The network thus develops as users of the network choose. But end-to-end also limits in certain respects how the network can develop. By limiting the functions at the network level, the design assures neutrality in how the network will develop.

Innovators need not fear that powerful actors within the network will bias the network against their developments. If users of the network demand it, then the network will provide it.

The end-to-end design also mirrors the values implicit in the common law structure of “common carrier” regulation. Like states with respect to interstate commerce, a common carrier must remain neutral about the service it provides. A common carrier must take all comers if it takes any. This neutrality assures entrepreneurs that they will not be vulnerable to at least this dimension of strategic cost. This in turn can lower the cost on innovation.[LEMLEY99]

4.2 Open Access

While the end-to-end Internet model was first adopted for technical reasons, it has important policy consequences as well. [LEMLEY99] The plug-and-play nature of Internet interoperability enables a wide variety of applications to connect; it therefore is architected to maximize the number of entities that can compete for the use of the network. [LEMLEY99] The Internet supports a complex and dynamic industry structure – one that offers market opportunities for many different companies. Although the overall industry structure can be quite hierarchical (because there is opportunity for specialized companies of many different sizes), individual management of each company can be flat and therefore able to rapidly respond to changing market conditions. [SALTZER99] By keeping the network simple and its interactions neutral, the Internet has facilitated applications that could not have been envisioned.[LEMLEY]

One consequence of the end-to-end design is that the network weakens the opportunity for strategic behavior by particular actors on the network, and by owners of the network. If control is vested in the ends, then choke points on the network are eliminated. Network owners cannot control how the network will develop. The advent of layer-violating network devices, however, makes interoperability problematic in a growing number of cases. Internet Service Providers (ISPs) would like to bundle access

to the Internet with a collection of other services such as Email, web-page hosting, storage, etc. For many customers this bundling is convenient. But technologically bundling these services can give a dominant ISP an anti-competitive motivation to use LV network devices to deny customers open access to the Internet for competing services.

Here we report five examples of potentially anticompetitive behavior enabled by deviating from end-to-end design: (summarized from [SALTZER99])

- 1) Fixed Backbone Choice. ISPs connect with international backbone networks (similar to long distance carriers in the telephone analogy) based on economic incentives, availability, and often cross-ownership. Besides a potential conflict-of-interest, an ISP backbone may prevent users from getting better service from another backbone.
- 2) Filtering. Several ISPs have begun to examine packets that they carry and discard those with certain purposes. Again there is a conflict-of-interest in that the ISP has an incentive to find technical or political logic to filter out competing services.
- 3) No Home Networks. In refusing to attach home networks, ISPs are actually protecting their ability to assign the IP address of the customer. By refusing to carry traffic to IP addresses they did not assign, the access provider can prevent the customer from contracting for a competing service from another ISP.
- 4) Server restrictions. Some ISPs impose an “acceptable use policy” that forbids customers from operating an Internet service such as a web site. The technical excuse is that web sites attract traffic and the provider has limited capacity. However, again the ISP has a conflict-of-interest because it offers a web site hosting service.

- 5) Content Limits. Some ISPs either limit the number the number of minutes or outright deny customers the use of “streaming video/audio” or download of MP3 audio (i.e., NAPSTER). The technical excuse for these restrictions is capacity constraint but the ISP has a conflict-of-interest, they will restrict new services that may sometime in the future become a competing services.

These five examples conflict with the Internet value of transparency. Saltzer reinterprets his end-to-end argument in this policy context:

“The end-to-end argument 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 features itself.”[SALTZER99]

4.3 Risks

No one can accurately predict the effect of increasing deployment of LV devices on the end-to-end Internet model. Our aim in this paper is not to claim that these threats will be realized. Our objective instead is simply to identify some of these potential costs. To the extent these costs exist, they evince the externality created by compromising on the end-to-end design. These potential risks include the following:

- 1) ISP market concentration: Compromising on end-to-end would increase the ability of content providers to discriminate in the content they offer. The ability to facilitate discrimination in the access to content across the Internet may facilitate concentration in what is currently a highly competitive ISP market. If service providers can guarantee premium access to certain “channels,” bundled with a dominant form of access, this can increase market concentration.
- 2) Control Innovation: By gaining control over Internet infrastructure, traditional companies are in a position to protect existing markets from a threat the Internet might create. For example, to the extent broadband service on the Internet might

be a threat to the existing market for cable services, cable companies would have an interest in protecting cable services from a broadband threat. To date, the behavior of some cable providers is consistent with the hypothesis that they would architect the network to protect their legacy monopoly. The MediaOne Internet service systems restrict customers to 10 minutes of video-streaming content at a time. When asked whether the company would permit streaming generally, an AT&T executive responded that AT&T didn't spend \$56 billion to get into the cable business "to have the blood sucked out of our vein." [USA 1999].

- 3) Threat to Innovation: To the extent that any actor can intervene to protect an existing technology, by blocking competing technologies, or by discriminating against them, this will increase the costs to innovators within that market. The threat of strategic action against new innovation will weaken the incentive for such innovation. (This, for example, was the theory of the government's case against Microsoft corporation.)

Compromising end-to-end can increase the risk of such strategic action. If the network can be architected to embed and protect one form of content distribution, for example, that will stifle innovation in other forms of distribution.

In all three ways, compromising on the principle of end-to-end presents risks to competition, and hence innovation, on the network. The change thus affects interests beyond the narrow interests at stake when any particular decision to implement a LV device is made. Some of the costs, in other words, of a LV device are born by the network. These costs, like pollution, are externalities to the network.

4.4 Governance

Our aim so far has been to suggest that compromising end-to-end as a principle of network design may have effects upon innovation and efficiency. It may also affect Internet governance. To the extent changes in network design concentrate power in

network owners, those changes will increase the power of individual actors to determine the evolution of network design. Such changes will thus increase the “governance” of the net, but not through institutions of traditional governance. In the world of perfect end-to-end design, no single institution could exercise power over how the network would develop. In a world where end-to-end is compromised, the potential for large institutions, controlling the use of aspects of the network, increases.

The compromise of end-to-end may also induce the emergence of institutions of traditional governance – ICANN in particular. This is because of the costs that LV devices impose on the existing Internet. LV devices increase the coordination costs for deploying new technologies to the existing Internet. With every new LV device, new applications must take account of the standards for those devices before they can be certain to run on the Internet generally. These coordination costs may be reduced if the network were to move to an IPv6 standard. The slowness of existing systems to move to the IPv6 standard, however, may induce governments to support organizations (such as ICANN) in their efforts to advance the network standard. This in turn may increase the jurisdiction of these international “governance” bodies, as they work to secure international network standards.

While the emergence of solutions to the costs imposed by LV devices through governance organizations may well be an improvement, one must also account for the risks that any governance organization will present beyond the specifics of this one case. Governance structures concentrate power in ways that may well stifle innovation and liberty on the net. These unintended consequences are another potential cost of deviating from the original network standard. [RESNICK 99]

5.0 Summary

In this paper we have presented the evolving challenges to the overall transparency of the end-to-end Internet model as proposed by Saltzer, Reed, and Clark. It can be argued that the transparency inherent to the end-to-end model is in many ways responsible for

the engineering success of the Internet. However, with unprecedented growth of the Internet has come pressure to violate the end-to-end model. We have documented examples of where Internet protocols designed for end-to-end transparency will not work in a world where packets have to traverse LV network devices such as NATs and firewalls. The large investment in LV network devices has been for valid reasons and this installed infrastructure will not be easily changed.

The trend continues toward incorporating more processing within the network. Active network research ranges from packets programming routers to routers making pre-programmed decisions based on packet content.¹⁵ [TENNENHOUSE97] In response to recent denial-of-service attacks, the IETF is convening an itrace BOF¹⁶ to process reverse path state information on packets within intermediate routers using ICMP traceback mechanisms. There is a tension building between providing end-systems knowledge of network conditions to provide enhanced services versus increased security vulnerabilities based on this knowledge.¹⁷ An example of this is the ongoing discussion within the IETF of an Internet draft on “Fog Lamps” to improve **visibility** of network devices to end-systems, a view directly opposed to the **transparency** view of the end-to-end Internet model.[LEAR00]

No one can predict the ultimate effects of layer violations on the Internet.[LESSIG99] In one scenario a complete migration to IPv6 potentially allows the restoration of a global address space and end-to-end transparency albeit with firewalls and PEPs still remaining. At the other extreme, only a partial IPv6 deployment leads to fragmentation of the network layer, with global connectivity resembling islands of connectivity. The Internet architecture has helped fuel the greatest economic boom in recent history; we should be

¹⁵ While some researchers have attacked “active networking” as not scalable and destabilizing, it is premature to make such determinations.

¹⁶ ICMP Traceback (itrace) Birds-of-a-Feather (BOF) at the 47th IETF meeting, Chair: Steve Bellovin, 3/30/00 15:30-17:30

¹⁷ Increased knowledge of conditions within the network may make additional diagnostic information available to interloping devices.

skeptical of changes in its design.[LESSIG99, LEMLEY99] The strong presumption should be in favor preserving the architectural features that have produced this extraordinary innovation.[LEMLEY99]

8.0 References

[ALLMAN99] Allman, M. et. al., “*Enhancing TCP Over Satellite Channels Using Standard Mechanisms*” RFC 2488, 1999.

[CLARK95] Clark, David D. “*The Design Philosophy of the DARPA Internet Protocols*” ACM Computer Communications Review, Vol. 25 No. 1, Jan. 1995, pp. 102-111.

[BLAKE98] Blake, S. et. al., “*An Architecture for Differentiated Service*”, RFC 2475, Dec. 1998.

[BRADEN97] Braden, R. et. al., “*Resource ReSerVation Protocol -- Version 1 Functional Specification*” RFC2205, Sept. 1997.

[CARPENTER00] Carpenter, Brian. “*Internet Transparency*” RFC 2775, 2000.

[CHEN98] Chen, Thomas M. and Alden Jackson et. al., “*Commentaries on Active Networking and End-to-End Arguments*” IEEE Network Magazine, May/June 1998.

[DEERING00] Deering, Steve, personal communications, Cisco Corp., March 2000.

[KENT98] Kent, S. and R. Atkinson. “*Security Architecture for the Internet Protocol*” RFC2401.

[KRUSE00] Kruse, Hans. “*The Pitfalls of Distributed Protocol Development: Unintentional Interactions between Network Operations and Applications Protocols.*” 8th Intl. Conf. on Telecom. Systems (ICTS), Nashville TN. USA, March 2000, pp. 289-293.

[KRUSE99] Kruse, Hans. “*Protocol Interactions and Their Effects on Internet-Based E-Commerce*” 2nd Intl. Conf. Telecom. and E-Commerce (ICTEC), Nashville. TN USA, Oct. 1999.

[HAIN00] Hain, T. “*Architectural Implications of NAT*” (work in progress)

[HOLDREGE00] Holdrege M and P. Srisuresh. “*Protocol Complications with the IP Network Address Translator (NAT)*” (work in progress, March 2000)

<http://www.ietf.org/internet-drafts/draft-ietf-nat-protocol-complications-02.txt>

[FLOYD00] Floyd, S. et. al., “*IPsec Interactions with ECN*” (work in progress)

<http://www.ietf.org/internet-drafts/draft-ipsec-ecn-00.txt>

[LEAR00] Lear, Eliot. “*NAT and other Network ‘Intelligence’: Clearing Architectural Haze through the Use of Fog Lamps*” (work in progress, December 1999)

<http://www.ietf.org/internet-drafts/draft-lear-foglamps-01.txt>

[LEMLEY99] Lemley, Mark A. and Lawrence Lessig. “*Ex Parte Comments in the Matter of Application for Consent to the Transfer of Control of Licenses MediaOne Group, Inc. to AT&T Corp.*”, FCC CS Docket No. 99-251.

<http://cyber.law.harvard.edu/works/lessig/cable/fcc/fcc.html>

[LESSIG99] Lessig, Lawrence. “*It’s the Architecture, Mr. Chairman.*”

<http://cyber.law.harvard.edu/works/lessig/cable/Cable.html>

[RESNICK99] Lessig, Lawrence, Resnick, Paul, *Zoning Speech on the Internet: A Legal and Technical Model*, 98 Mich. L. Rev. 395 (1999).

[SALTZER99] Saltzer, Jerome H. “ ‘Open Access’ is Just the Tip of the Iceberg” Oct. 22, 1999. <http://web.mit.edu/Saltzer/www/publications/openaccess.html>

[SALTZER84] Saltzer, Jerome H., Reed David P. and David D. Clark. “*End-to-End Arguments in System Design*” ACM Trans. on Comp. Systems, Vol. 2 No. 4, Nov. 1984, pp. 277-288. (an earlier version appeared in 2nd Intl. Conf. on Distr. Computer Systems, April 1981, pp. 509-512.)

[TENNENHOUSE97] Tennenhouse, D. et. al., “*A Survey of Active Research Network Research*” IEEE Communications Magazine, Vol. 35 No. 1, 1997.

[TERZIS00] Terzis A. “*RSVP Operation Over IP Tunnels*”, RFC 2746, 2000.
