The Privatization of the Internet's Backbone Network

Rajiv C. Shah & Jay P. Kesan

Rajiv C. Shah
Adjunct Assistant Professor
University of Illinois at Chicago
rshah@a5.com

Jay P. Kesan (Corresponding Author)
Professor
College of Law
University of Illinois at Urbana-Champaign
504 E. Pennsylvania
Champaign, IL 61820
T (217) 333-7887
kesan@law.uiuc.edu

This material is based upon work supported by the National Science Foundation under Grant No. ITR-0081426 & IIS- 0429217. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Abstract

Scholars have neglected the privatization of the Internet's backbone network, despite the obvious significance of the U.S. Government turning control over a powerful new communication technology to the private sector. This paper describes the transition from a government sponsored backbone network to multiple commercially owned backbone networks. We also analyze the implications of the privatization upon the Internet's governance, competition, and performance.

The Privatization of the Internet's Backbone Network

Histories of the Internet abound (Abbate, 1999; Hauben & Hauben, 1997; Kahn & Cerf, 1999; Moschovitis, Poole, Schuyler, & Senft, 1999; Naughton, 2000; Reid, 1997; Salus, 1995; Schiller, 1999; Segaller, 1998), yet a comprehensive account of the privatization of the Internet's backbone network does not exist. When it comes to describing the transition of control from the government to the private sector, the descriptions suddenly shift to the passive tense. It becomes unclear who the actors were and what actions they took to privatize the backbone network. The lack of coverage from the mainstream press led Project Censored to place the privatization of the Internet backbone in their top ten list for 1995 (Jensen, 1997).

This article provides the history of the privatization of the backbone network. This is valuable because the privatization has not been scrutinized by either academics or the press. We also believe the privatization has a significance for communication scholars for several additional reasons. First, it represents the transfer of a significant communication technology to the private sector. As a result, the private sector inherited a technology that was created with billions of public dollars (Lytel, 1998; MacKie-Mason & Varian, 1994). The transfer also meant a shift from government control to private control over network resources. Unlike other communication technologies, the privatization of the backbone left little regulatory requirements. As we later point out, private backbone networks are not subject to "must carry" regulations or nondiscriminatory standards for the traffic over their networks. Second, an historical understanding of the privatization provides key insights into contemporary issues, such as the performance of the Internet, competition in the backbone industry, and the governance of the Internet. Finally, this is not the only privatization involving the Internet. The government is now transitioning the Domain Name System (DNS) and the Internet Protocol (IP) address system, both key components of the Internet, to private control. Therefore, any lessons learned from the privatization of the backbone network may aid these other Internet privatizations.

The privatization of the backbone network involved reshaping the National Science Foundation Network (NSFNET) into what is known today as the Internet. This process affected both the content across the NSFNET as well as the control of the underlying infrastructure. The actual privatization consisted of government shifting from the practice of contracting out a government-subsidized backbone to allowing the market to provide backbone services. The first part of this paper discusses this privatization process. The second part focuses on several important implications of the privatization, which are relevant to ongoing and future privatizations. We also discuss how the government lost an opportunity to ensure societal values were considered in the design of the Internet.

The Privatization Process

The critical step in the creation of the Internet was the development of the NSFNET. This network was created because of an interest in supercomputing in the 1980s (Rogers, 1998). The very nature of supercomputing only allowed the creation of a few facilities. To ensure the research community had access to the supercomputing centers, the National Science Foundation (NSF) set up a national research network in two stages. The first stage involved interconnecting five NSF funded supercomputer centers. The second stage involved constructing a high-speed backbone to link together various regional and campus networks across the United States.

By 1986, the NSFNET was operational. The NSFNET backbone connected more than 200 colleges and universities together via the regional networks (Office of Technology Assessment, 1993). The regional networks connected to smaller local networks, such as universities. The NSFNET also connected to various federal networks, such as the Department of

Energy's Energy Sciences Network and NASA's NASA Science Internet. The connection of all these networks via the NSFNET was known as the Internet. The structure and role of the NSFNET in connecting the regional networks and the local networks is shown in Figure 1.

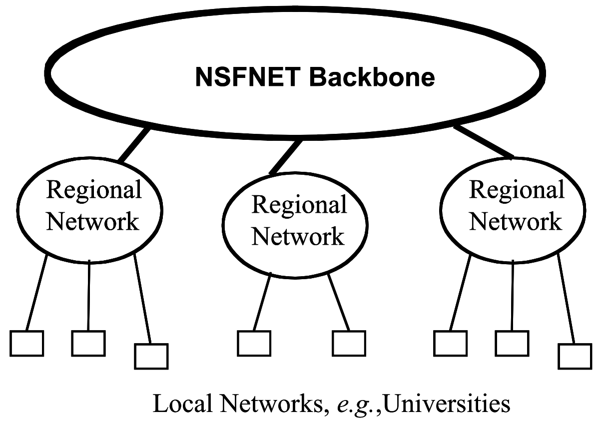

Figure 1. Architecture of the NSFNET

This research network was supported by funds from the NSF, state, university, federal, and private sector (*Computer Networks and High Performance Computing*, 1988, 66-67). Besides offering grants to encourage academic institutions to connect to the NSFNET via the regional networks, the NSF also encouraged the regional networks to find commercial customers. The NSF reasoned that the revenue from new customers would allow regional networks to expand and use the economies of scale to lower costs for everyone (Leiner et al., 1997).

While this network was adequate for educational uses, there was a movement to privatize the NSFNET. There was little public debate or opposition to the privatization. By the early 1990s, telecommunications policy for both political parties was based upon notions of deregulation and competition (Olufs, 1999). At numerous junctures before the privatization of the NSFNET, politicians, and telecommunication executives made it clear that the private sector would own and operate the Internet. In fact, the privatization of the Internet became seen as virtually inevitable. This inevitability was spurred on for two reasons. The first was the desire to send private, non-governmental traffic across the Internet. The second was the wish to sell connectivity or infrastructure by telecommunications companies.

Acceptable Use Policy (AUP)

The NSFNET connected universities, federal agencies, public and private research laboratories, and community networks. While the NSF encouraged such diversity, it also had an Acceptable Use Policy (AUP). The AUP prohibited the use of the NSFNET for purposes not in support of research and education, a policy consistent with the NSF's mission (Kahin, 1992). Nevertheless, a growing number of users wished to use NSFNET for purposes beyond research and education, a push for what the NSF termed "commercial use" (Kahin, 1990). The first official commercial use of the NSFNET began in 1988 with Internet pioneer Vinton Cerf. He persuaded the government to allow MCI Mail to link with federal networks—the Internet—for "experimental use." Soon after, other companies such as CompuServe and Sprint gained permission for "experimental use" for commercial email. The stated rationale was that these commercial providers would enhance research and educational uses by allowing researchers to communicate with more people (Adam, 1996).

The NSF recognized that by disallowing commercial traffic, it would encourage the growth of commercial backbone companies to create AUP free networks (Leiner et al., 1997). This strategy worked. Virtually all of the later commercial backbone providers emerged as for-profit spin-offs from the nonprofit regional networks (Cerf, 2000). For example, in late 1989, two of the founders of the nonprofit New York regional network, NSYERNet, established a for-profit company, Performance Systems International (PSI). PSI would become a backbone provider and has recently been acquired by Cogent Communications.

These new firms provided Internet connectivity to companies who were fearful of violating the NSF's AUP. Eventually, several of these commercial backbone companies decided to interconnect their networks, creating the Commercial Internet Exchange (CIX) (Chinoy & Salo, 1997). By creating the CIX, the commercial backbones could access users on their competitors' backbones. Moreover, the commercial backbones could gain these additional users without violating the NSF's AUP. In effect, the commercial backbone providers created a commercial alternative to the NSFNET.

The NSF contracted out the management and the operation of the NSFNET backbone to the team of Michigan Educational Research Information Triad (MERIT), MCI, and IBM. The emerging opportunities for commercial backbone companies led the NSFNET's contractors to enter this new industry. In September 1990, this team created a not-for-profit corporation, Advanced Network Services (ANS), as a subcontractor to operate the NSFNET backbone ("MCI and IBM Form Nonprofit Supercomputing Company," 1990). ANS would also offer networking services on the new network ANSNet. This new network shared the same infrastructure as the NSFNET. This allowed ANS to provide commercial connectivity to the NSFNET, unlike its other commercial competitors (1991).

The relationship between ANS and the NSF drew criticism and charges of unfair competition by other commercial backbone providers (Markoff, 1991). In response to the criticism, congressional hearings were held on the management of the NSFNET (*Management of NSFNET*, 1992). As a result of the hearings, the Office of the Inspector General (OIG) performed a review of the NSFNET's management (1993), and discovered a number of problems (Kesan & Shah, 2001). One response to these problems was a bill to remove the NSF's AUP and allow commercial traffic across the NSFNET. This bill was amended later to allow commercial use of the network as long as it would increase the networks' utility for research and education ("Scientific and Advanced-Technology Act," 1992). In late 1992, the NSF's new AUP allowed the private sector to use the network as long as it indirectly benefited research and education. This new, liberal AUP led to even more growth for the NSFNET as commercial firms could use the NSFNET for a variety of reasons (Johnson, 1993).

Infrastructure

The privatization of the NSFNET infrastructure had been expected for a long time. The prominent computer scientist, Leonard Kleinrock, representing the National Research Council on networking, told Congress in 1988 that any government-run network would eventually be transitioned to the telecommunications industry (*Computer Networks and High Performance Computing*, 1988). Similarly, a bill introduced in Congress in 1989 acknowledged eventual control and ownership of the infrastructure by commercial providers. The language of this bill by Senator Al Gore highlighted the importance of the private sector. In regards to a national network, the bill stated that the network shall be privatized (*National High Performance Computer Technology Act*, 1989).

In 1989, the White House's Office of Science and Technology Policy (OSTP) produced a report with a detailed program plan for implementing the national network (Office of Science and Technology Policy, 1989). The plan called for a three-stage development process for the network, which would become the framework for privatizing the NSFNET backbone. The first stage, which the report noted was already underway, was to upgrade networks to T1 (1.5 Mbs/sec) speed. In the second stage, the backbone was to be upgraded to T3 (45 Mbs/sec) speed. The report also stated that stage two, "will provide a base from which commercial providers can offer compatible networking services nationally," (Office of Science and Technology Policy, 1989, 35). The third stage would involve a national backbone of 1 gigabit per second for a few sites. All other sites would connect at T3 speed. The report noted that the stage three deployment was not expected until the middle or late 1990s. The report went on to state that “the deployment of the Stage 3 NREN will include a specific, structured process resulting in transition of the network from a government operation to a commercial service,” (Office of Science and Technology Policy, 1989, 35). Because the plan did not detail exactly how the NSFNET was to be privatized, the NSF would proceed to fund studies on how to privatize the Internet through 1990 and 1991.

By late 1991, after these workshops and meetings, the NSF formulated a new design for the NSFNET with multiple backbones. The NSF announced this design in its Project Development Plan, which provided an overview for the future of NSFNET (National Science Foundation, 1992). The goal of this plan was to award competitive backbone services by April 1993 (Wolff, 1992). In response, the NSF received more than 240 pages of comments from industry groups. The comments largely expressed concern that only a few firms would be awarded the entire contract and, suggested instead, the desire for a more competitive network design (Anderson, 1993).

In the spring of 1993, the NSF incorporated the industry groups’ comments and released a revised solicitation with three parts (National Science Foundation, 1993). The first part consisted of a Routing Arbiter which operates as a “traffic cop” to ensure consistent routing policies (Lawler, 1995). The second part proposed the creation of a Very High Speed Backbone Service (vBNS) to replace the NSFNET as a new high-speed backbone for research and educational use. The vBNS would be initially restricted to a select group of researchers requiring high-speed networking for specialized applications. The third part of the solicitation concerned the use of Network Access Points (NAPs) to connect together the vBNS, federal networks, and commercial backbone networks. The concept of multiple backbones interconnecting through NAPs is shown in Figure 2. Instead of a central backbone connecting the regional networks, the regional networks had to choose a commercial backbone network. To ensure connectivity between commercial backbone networks, they would interconnect at the NAPs.

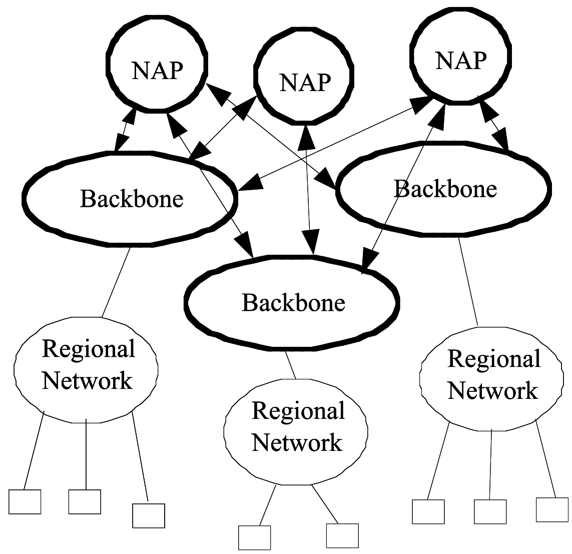

Figure 2. The New Design for the NSFNET

In February 1994, the NSF announced the award of the NAPs to Sprint for New York, MFS for Washington, D.C., Ameritech for Chicago, and Pacific Bell for California. NSF awarded MCI the contract to operate and maintain the vBNS. The regional networks were supposed to disconnect from the NSFNET backbone by October of 1994 and connect to commercial providers, which would interconnect via the NAPs. None of the regional networks were able to meet the October deadline, but over the following six months they migrated to commercial providers (Fazio, 1995). Finally, on April 30, 1995, the NSFNET was retired.

In August 1996, the NSF announced it would end its sponsorship of the four public NAPs. The NAPs would be operated entirely by the private sector (Farnon & Huddle, 1997). By this point, the government had transitioned from contracting out services to allowing the market to provide Internet backbone service—thus fully privatizing the Internet backbone network. After this decision, the only remaining NSF-subsidized backbone service was the vBNS. The vBNS served as the initial foundation of the ongoing Internet2 networking effort.

Implications of the Privatization

The privatization of the NSFNET backbone fundamentally reshaped the infrastructure of the Internet. Instead of one major backbone—the NSFNET—the new network depended on multiple backbone providers. The NSF created the NAPs to interconnect these networks and to prevent a balkanized Internet. This new network was designed to create competition for backbone services. In designing this new network, the NSF put little constraints on the design and use of the network. Specifically, the NSF did not set forth any policies or requirements on the future of the new networks.

This section discusses how this hands-off approach has shaped the Internet with predictable implications. First, the lack of any performance requirements for the public NAPs has led to them becoming congested. This has led to new technological innovations that are affecting how information is transferred across the Internet. Second, the lack of an interconnection policy has resulted in a concentrated industry with a few large backbone providers. As a result, there is a high entry barrier to the backbone market along with concerns that a few backbone providers may act anti-competitively. Third, the government's hands-off regulatory policy ensures the private industry governs and controls the Internet's backbone. A final implication concerns the lost opportunity by government to ensure societal values were considered in the design of the Internet.

Performance of the Network Access Points

The lack of any performance requirements for the public NAPs has led to them becoming congested. In designing the NAPs, the NSF neglected to put into place minimum performance requirements. As a result, they soon became congested. For example, Sprint's backbone network operates at 2.5 billion bits per second, while the highest speed offered at the Sprint managed public NAP is 45 million bits per second. This is aptly characterized as "giving drivers on a six-lane highway access via a dirt road," (Weinberg, 2000, 238). This has led many backbone providers to consider the public NAPs' connections as worthless (Cook, 1998).

The movement away from the public NAPs has led to the use of private exchange points. Private exchanges points were contemplated as part of the NSF's redesign of the Internet. However, their recent popularity is a result of the congestion at the public NAPs (Winkleman, 1998). By interconnecting at private exchange points, it is possible to avoid routing traffic through the public NAPs. For example, the five large backbone providers have each implemented at least four connections with the other large backbone providers at private exchanges. This allows the large backbone providers to provide high performance Internet access for their customers. Meanwhile, smaller networks are often relegated to the congested public NAPs because there is no requirement for a private exchange to treat other networks in a fair and nondiscriminatory manner (Caruso, 2000).

The congestion at the NAPs has also led to the use of new technologies, which fundamentally affect how information is transferred across the Internet. The first is the Content Delivery Networks (CDNs) such as Akamai. The CDNs attempt to bypass as much of the public Internet as possible. Instead, of having one server for a web site's content, a CDN uses a number of servers with replicated content at the edge of the network. Users are then transparently routed to the nearest available server. This allows for content to be delivered much more quickly to users (Phifer, 1999). This service is not free. As a result, web sites that can afford to pay for a CDN will be able to provide richer content much more quickly than other web sites (Sandvig, 2003; Spinrad, 1999).

A second technology used to address congestion on the Internet is known as Quality of Service (Ferguson & Huston, 1998). This technology modifies one of the basic tenets of Internet traffic. Traditionally, all Internet traffic was treated alike and everyone would give their "best effort" to transmit all traffic. However, new technologies are being deployed to provide preferential treatment to traffic. This preferential treatment has obvious technical applications. For example, networks may make a high-speed lane for video conferencing and a slow lane for email traffic. One popular recent use of Quality of Service is by universities seeking to limit the bandwidth of file sharing applications on their networks (Bray, 2002).

The ability to segregate traffic by application, users, or network can has enormous implication for society's use of the Internet. Quality of Service now allows service providers to differentially treat their customers on the basis on who they are or what they are doing. One controversial use of Quality of Service is by cable companies providing Internet access. Quality of service technologies could allow them to manage network traffic to expedite the delivery of affiliated content while demoting competitive material to second-class service (Lemley & Lessig, 2001). Additionally, the ability to segregate traffic also impacts the backbone industry by favoring large backbone providers over smaller ones. For effective Quality of Service across a network, it is necessary to maintain control of the network traffic over the entire length of the network. Under the current system, smaller networks depend upon a few large backbone providers for some of their Internet transport. This limits their use of Quality of Service and provides the large backbone providers with a technological edge.

Competition and the Lack of an Interconnection Policy

The lack of an interconnection policy places no constraints on backbone providers. The current lack of a policy favors a few large backbone providers to the detriment of consumers and smaller backbone providers. At its simplest, an interconnection policy would ensure all competitors fair and nondiscriminatory access to each other's networks. An interconnection policy is important because to connect to the Internet, a network must connect with other networks.

The government has historically used interconnection policies to require entities to maintain a duty to serve and interconnect on a nondiscriminatory basis. For example, steamships, railroads, and the telephone network were all subject to policies to ensure interconnection, which served to foster and grow their networks (Speta, 2002). Similarly, the Telecommunications Act of 1996 requires telecommunications carriers to interconnect. Once again, this is because Congress implicitly understood that a telecommunication network's value increased as more people became connected to the network. Therefore, it is in the public interest for networks to be able to effectively interconnect. While these laws apply to traditional voice telecommunication carriers, the Federal Communications Commission (FCC) has held that Internet backbone providers are not telecommunications carriers as defined in the 1934 or 1996 Telecommunications Acts (Kende, 2000). Thus, there is no current interconnection policy required by law for backbone services.

Without an interconnection policy, the large backbone providers decide with whom they wish to peer and the conditions of the peering (Frieden, 1998). First, large backbone providers do not publicly disclose with what other networks they peer or their terms for the peering. This secrecy benefits the large backbone providers as they undoubtedly claim they are well-connected (Angel, 2000; Gareiss, 1999). Second, smaller networks often have to sign nondisclosure agreements, which limits sharing information with other smaller networks regarding their interconnection terms (Cukier, 1997). This practice has also limited the creation of industry standards concerning who should get free peering and what are reasonable rates (Weinberg, 2000). Third, larger backbone providers can discontinue their interconnection arrangements with little notice (Cukier, 1997). Such action disparately affects the smaller networks.

The lack of an interconnection policy has led to increased barriers to entry in the backbone industry. In 1995, there were five major backbone providers: UUNET, ANS, SprintLink, BBN, and MCI (Winkleman, 1998). Despite the rapid growth in Internet services and new Internet start-ups such as Level 3 Communications and Qwest, the same companies still control the Internet. In 2000, there were five major backbone providers: MCI WorldCom (bought ANS and UUNET), Sprint, Genuity (formerly GTE who bought BBN), AT&T, and Cable & Wireless (which owns MCI's old backbone) (Weinberg, 2000). AT&T built its network with its own large fiber optic network and by acquiring IBM's Global Network and one of the early Internet backbone providers, CERFnet. After all, the only way to become a large backbone provider is to buy one!

And a few years later, nothing has changed except that its cheaper to buy a backbone provider with the MCI WorldCom in bankruptcy and Genuity in financial difficulty. Despite the financial problems, MCI WorldCom is still by far the largest backbone provider, according to TeleGeography Research. Moreover, Qwest and Level 3 Communications, both of which have considerable excess capacity are also in financial trouble. As a result, AT&T and Sprint will likely increase their market share as firms seek financially stable providers (Alster, 2002). A final wildcard that could affect the backbone industry is the entry of the financially healthy baby bells into the backbone provider market. Regulators have long barred them from offering backbone services, because of the lack of competition in their local markets.

Regulation of Internet Backbone Providers

In today's privatized Internet, the private sector governs the Internet. The FCC has maintained a hands-off approach with the Internet backbone providers. This lenient regulatory approach began in the early 1980s when the FCC ruled that enhanced services were not subject to common carrier regulation (Oxman, 1999). This removed Internet services from the many regulations that govern voice telephony. In contrast, ISPs benefit from FCC regulated interconnection policies with local phone companies. The only scrutiny of the backbone industry came with the merger of MCI and WorldCom in 1998 (Schiller, 1998). Regulators, led by the Justice Department, ensured that MCI sold off its Internet assets before the merger to prevent MCI WorldCom from owning 70% of the Internet (Goodin, 1998).

Although the Internet is not regulated, it has been able to maintain its robustness through the years. The Internet is governed by Scott Bradner's belief that "enlightened self-interest among the many different groups that make up the Internet will drive cooperation in the future," (Cooney, Gaffin, & Messmer, 1995, 1). The prevailing wisdom as represented by David Farber, former Chief Technologist at the FCC, is that these companies have a self-interest in the success of the Internet (Caruso, 2000). The use of self-interest as the primary motivating factor for management of the Internet has worked well in maintaining system performance and reliability. For example, the Internet was able to easily withstand the shutdown of the Ebone network owned by KPNQwest. Traffic was routed onto other networks, which resulted in slight delays on the Internet traffic in Europe, but no major problems (Evers, 2002; Hansell, 2002).

One of the consequences of the lack of regulatory oversight is that government cannot provide consumers with assistance when something goes wrong. For example, the FCC Chairman Michael Powell acknowledges that he cannot stop backbone providers from shutting down their service, while he could prevent phone service from being stopped. Moreover, backbone providers could even legally ignore FCC directives and cut off service to customers with little notice, thus leaving them little or no time to switch providers (Dreazen, 2002). This type of rapid shutdown has occurred. When the backbone provider PSI failed to meets its obligations with Cable & Wireless it was unilaterally terminated. This resulted in customers unexpectedly without Internet service (Kraph, 2001). Similarly, when Excite@Home went bankrupt it left nearly a million cable modem users without Internet access for a week (Konrad, 2001). These types of interruptions would not be permitted for phone service due to stricter regulatory oversight.

Merely because the Internet industry is governed by self-interest does not mean that government has no interest in overseeing the Internet. Potential conflicts of interest exist between large backbone providers and their competitors, customers, or some public interest. This issue is especially important when it comes to network standards and new networking technologies. For example, large backbone providers could use technologies that adversely affect the public interest, for example privacy or competition. In these cases government would need to step in if the market could not address these issues. Second, the large backbone providers have enormous power in introducing new technologies. Decisions to introduce new technologies such as next generation transmission standard, IPv6, are in their hands. If they do not implement it, nobody will. And if they require everyone else to implement it, everyone else will have to. It is also important to note that the large backbone providers operate without any significant checks on their behavior. For example, they already meet regularly in private to discuss engineering issues (Cukier, 1997). Thus, a few large backbone providers have the ability to use new technologies in a manner that could be detrimental to the public interest and competition. Government must maintain vigilance to ensure that the large backbone providers do not manipulate technologies merely for their own benefit.

Lost Opportunity by the NSF

One final implication of the privatization was the NSF's lost opportunity to address societal concerns when redesigning the Internet's technological infrastructure. For example, the NSF could have addressed the lack of security on the NSFNET. Security issues with the NSFNET were recognized back in 1987 (Office of Science and Technology Policy, 1987). However, over the NSFNET's lifetime, the government never seriously addressed concerns of security. The most significant development for Internet security was the creation of the Computer Emergency Response Team (CERT). CERT is a reporting center for Internet security problems that disseminates security information to the Internet community. While CERT helps people who are affected by security issues, it does not perform fundamental research and development to address security issues.

The significant point is that government could have done something, not necessarily that the government should have. However, there was little public debate concerning what the government could have done. The lack of government action is especially appalling when you consider the potential influence it may have had. The government lost a valuable opportunity during the redesign of the Internet to ensure that social concerns were addressed in the redesigned Internet. For example, the Internet could have been designed in a more secure fashion by requiring security mechanisms. In fact, there is still a movement to add more security to the Internet by using more secure network protocols. For example, the next generation Internet Protocol, IPv6, will increase security (Hagen, 2002). In sum, the government failed to consider societal concerns and, as a result, lost its onetime opportunity during the redesign of the Internet.

Conclusion

This paper describes the privatization process for the backbone network as well some of the implications of the privatization. Our telling of the history identified three important facets of the privatization. First, we identified several reasons for the privatization. These include privatization as the dominant policy approach, the desire to send commercial content across the Internet, and allowing multiple telecommunication companies to own and control the infrastructure of the Internet. Second, we found that the primary actor in this process was the NSF, which was following the dictates of Congress and the Executive branch. Finally, the privatization was manifested by the NSF removing its central backbone and funding the NAPs to interconnect commercial networks.

We also identified four important implications of the privatizations, which include performance issues with the Internet, competition in the backbone industry, the governance of the Internet, and lost opportunities. We believe that these implications are also relevant to the other privatizations of the Internet consisting of the Domain Name System (DNS) and the allocation of IP addresses. It is well established that the privatization of the DNS is already suffering from problems with competition and governance (Froomkin, 2000; Kesan & Shah, 2001; Mueller, 2002). Our hope is that our work will aid in preventing these problems from being repeated in future privatizations.

References

Abbate, J. (1999). *Inventing the Internet*. Cambridge, MA: MIT Press.

Adam, J. (1996, November). Geek Gods: How Cybergeniuses Bob Kahn and Vint Cerf Turned a Pentagon Project into the Internet and Connected the World. *Washingtonian, 32,* 66-78.

Alster, N. (2002, Sep. 1). Telecom Wreck Provides an Opening for AT&T. *New York Times,* p. 3.4.

Anderson, C. (1993). The Rocky Road to a Data Highway. *Science, 260*(5111), 1064-1065.

Angel, J. (2000, Feb. 1). Toll Lanes on the Information Superhighway. *Network Magazine, 15,* 42-49.

Bray, H. (2002, Oct. 7). Getting Tough Against College' File-Swappers. *Boston Globe,* p. C.3.

Caruso, D. (2000, February 14). The Internet Relies on Networks' Passing Data to One Another. But What Happens if One of Them Refuses? *New York Times,* p. C.4.

Cerf, V. G. (2000). *Computer Networking: Global Infrastructure for the 21st Century*. Retrieved September 10, 2003, from http://www.cs.washington.edu/homes/lazowska/cra/networks.html

Chinoy, B., & Salo, T. J. (1997). Internet Exchanges: Policy-Drive Evolution. In B. Kahin & J. H. Keller (Eds.), *Coordinating the Internet* (pp. 325-345). Cambridge: The MIT Press.

Computer Networks and High Performance Computing, US Senate, Second Sess. (1988).

Cook, G. (1998). *July COOK Report on the Internet*: www.cookreport.com.

Cooney, M., Gaffin, A., & Messmer, E. (1995, April 3). Internet Surge Strains Already Shaky Structure. *Network World, 12,* 1, 67.

Cukier, K. N. (1997). *Peering and Fearing: ISP Interconnection and Regulatory Issues.* Paper presented at the Impact of the Internet on Communications Policy, Cambridge, Massachusetts.

Dreazen, Y. (2002, July 31). FCC Chief Cites Internet Weak Spot*. Wall Street Journal,* p. A2.

Evers, J. (2002, July 3). No Internet Delays After Partial KPNQwest Shutdown*. IDG News Service*.

Farnon, M., & Huddle, S. (1997). Settlement Systems for the Internet. In B. Kahin & J. H. Keller (Eds.), *Coordinating the Internet* (pp. 377-403). Cambridge: The MIT Press.

Fazio, D. (1995, March 14). *Hang onto your Packets: The Information Superhighway heads to Valleyfair*. Retrieved Sep. 10, 2003, from http://ei.cs.vt.edu/~history/internet.hist.html

Ferguson, P., & Huston, G. (1998). *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*. New York: John Wiley & Sons, Inc.

Frieden, R. (1998). Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. *Virginia Journal of Law and Technology, 3*, 8.

Froomkin, M. (2000). Using ICANN to Route Around the APA and the Constitution. *Duke Law Journal, 50*, 17-186.

Gareiss, R. (1999, October 7). Old Boys' Network. *Data Communications, 28,* 36-46.

Goodin, D. (1998, Mar. 16). WorldCom-MCI Scrutiny Heats Up. *CNET News.com*.

Hagen, S. (2002). *IPv6 Essentials*. Sebastopol, CA: O'Reilly & Associates.

Hansell, S. (2002, July 24). Among WorldCom Puzzles, Future of UUNet Service*. New York Times,* p. C.8.

Hauben, M., & Hauben, R. (1997). *Netizens: On the History and Impact of the Usenet and the Internet*. Los Alamitos, CA.

Jensen, C. (1997). *20 Years of Censored News*. New York: Seven Stories Press.

Johnson, J. T. (1993, Mar.). The Internet Opens Us to Commercial Use. *Data Communications, 22,* 55-60.

Kahin, B. (1990). *RFC 1192: Commercialization of the Internet*. Retrieved September 10, 2003, from http://www.faqs.org/rfcs/rfc1192.html

Kahin, B. (1992). The NREN as Information Market: Dynamics of Public, Private, and Academic Publishing. In B. Kahin (Ed.), *Building Information Infrastructure* (pp. 323-401): McGraw-Hill Primis.

Kahn, R. E., & Cerf, V. G. (1999, December). *What is the Internet (And What Makes It Work)*. Retrieved September 10, 2003, from http://www.policyscience.net/cerf.pdf

Kende, M. (2000). *The Digital Handshake: Connecting Internet Backbones* (No. OPP Working Paper No. 32): Office of Plans and Policy, Federal Communications Commission.

Kesan, J. P., & Shah, R. C. (2001). Fool Us Once Shame on You -- Fool us Twice Shame on Us: What We Can Learn from the Privatizations of the Internet Backbone Network and the Domain Name System. *Washington University Law Quarterly, 79*(1), 89-220.

Konrad, R. (2001, December 26). Oversight Needed? Mishaps Fuel Regulation Debate. *CNET News.com*.

Kraph, E. (2001, July). Straightening Out the Internet Backbone. *Business Communication Review, 31,* 14-15.

Lawler, A. (1995, 17 March). NSF Hands Over the Internet . . . National Science Foundation to Privatize NSFNET. *Science, 267,* 1584.

Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., et al. (1997). The Past and Future History of the Internet; The Next 50 Years: Our Hopes, Our Visions, Our Plans. *Communications of the ACM, 40*(2), 102-108.

Lemley, M. A., & Lessig, L. (2001). The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era. *UCLA Law Review, 48*(4), 925-972.

Lytel, D. (1998, March). Nonprofit Parents, Corporate Kids. *Upside Today, 10,* 66.

MacKie-Mason, J. K., & Varian, H. (1994). The Economics of the Internet. *Dr. Dobb's Journal, 19*(14), 6-16.

Management of NSFNET, US House of Representatives (1992).

Markoff, J. (1991, December 19). Data Network Raises Monopoly Fear. *The New York Times,* p. D.7.

MCI and IBM Form Nonprofit Supercomputing Company. (1990, Sep 18). *Communications Daily, 10,* 3.

Moschovitis, C. J. P., Poole, H., Schuyler, T., & Senft, T. M. (1999). *History of the Internet: A Chronology 1843 to the Present*. Santa Barbara, CA: ABC-CLIO.

Mueller, M. L. (2002). *Ruling the Root: Internet Governance and the Taming of Cyberspace*. Cambridge: MIT Press.

National High Performance Computer Technology Act, US Senate (1989).

National Science Foundation. (1992, June). *Public Draft: Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services provider for NSFNET and the NREN program*. Retrieved September 10, 2003, from ftp://nic.merit.edu/nsfnet/recompete/solicit_comment.ascii

National Science Foundation. (1993, June). *Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for NSFNET and the NREN program*. Retrieved September 10, 2003, from ftp://nic.merit.edu/nsfnet/recompete/nsf9352

Naughton, J. (2000). *A Brief History of the Future: From Radio Days to Internet Years in a Lifetime*. Woodstock, NY: Overlook Press.

Office of Inspector General. (1993). *OIG Review of NSFNET* (No. OIG9301): National Science Foundation.

Office of Science and Technology Policy. (1987). *A Research and Development Strategy for High Performance Computing*. Washington, DC: Executive Office of the President.

Office of Science and Technology Policy. (1989). *The Federal High Performance Computing Program*. Washington, DC: Executive Office of the President.

Office of Technology Assessment. (1993). *Advanced Network Technology -- Background Paper* (No. OTA-BP-TCT-101). Washington, DC: U.S. Government Printing Office.

Olufs, D. (1999). *The Making of Telecommunications Policy*. Boulder, CO: Lynne Rienner Publishers.

Oxman, J. (1999). *The FCC and the Unregulation of the Internet* (No. OPP Working Paper No. 31). Washington, DC: Office of Plans and Policy, Federal Communications Commission.

Phifer, L. (1999, Nov. 30). Content Delivery Networks. *ISP Planet*.

Reid, R. H. (1997). *Architects of the Web: 1,000 Days that Built the Future of Business*. New York, NY: John Wiley & Sons, Inc.

Rogers, J. D. (1998). Internetworking and the Politics of Science: NSFNET in Internet History. *The Information Society, 14*(3), 213-228.

Salus, P. H. (1995). *Casting the Net: From ARPANET to Internet and Beyond*. Reading, MA: Addison-Wesley.

Sandvig, C. (2003, May). *The Structural Problems of the Internet for Cultural Policy.* Paper presented at the Critical Cyberculture Studies: Current Terrains, Future Directions, Seattle, Washington.

Schiller, D. (1998). *Bad Deal of the Century: The Worrisome Implications of the WorldCom-MCI Merger*. Washington, DC: Economic Policy Institute.

Schiller, D. (1999). *Digital Capitalism: Networking the Global Market System*. Cambridge, MA: MIT Press.

Scientific and Advanced-Technology Act, 42 U.S.C. § 1862(g) (1992).

Segaller, S. (1998). *Nerds 2.0.1: A Brief History of the Internet*. New York: TV Books.

Speta, J. B. (2002). A Common Carrier Approach to Internet Interconnection. *Federal Communications Law Journal, 54*(4), 225-280.

Spinrad, P. (1999, Aug.). The New Cool: Akamai Overcomes the Internet's Hot Spot Problem. *Wired, 7,* 152-158.

Valovic, T. (1991, July). ANS and Internet Commercialization. *Telecommunications Americas, 25,* 4.

Weinberg, N. (2000, June 12). Backbone Bullies. *Forbes,* 236-238.

Winkleman, J. (1998, August 15). Getting Connected. *America's Network, 102,* 30-34.

Wolff, S. (1992, November 26). *NSFNET Backbone Services after November 1992*. Retrieved March 31, 1999, from ftp://nic.merit.edu/cise/pdp.txt