The Evolution of Internet Interconnections

Sean Butler

Capstone Project

August 2000

M.S. Information Technologies

Rochester Institute of Technology

Table of Contents

Abstract

In 1995, the NSF officially shut down the NSFNet backbone, thereby ending the nascent Internets early architecture as a single backbone network. Today, the Internet is a group of loosely interconnected networks run by many diverse companies. These interconnections are in no way controlled by any industry or government agency, and are therefore held together only by the market demands of the Internet community. Although the FCC has traditionally maintained a stance of "unregulation" of all information and computer networks, they have increasingly show interest in ensuring the rapid deployment of Internet access. In addition, as more and more critical elements of communication are implemented on the Internet, some safeguards ensuring end-to-end connectivity, and therefore on maintaining the interconnection between networks, are needed. This paper discusses the history and evolution of Internet interconnections, compares and contrasts them to traditional telephony interconnections, and explores the possibility of regulation over such connections. This paper covers events up to the end of 1999.

Introduction

Today's Internet grew out of the National Science Foundation's (NSF) NSFNet, which was an extension of the experimental US Department of Defense's ARPANET. In 1993, after years of exponential growth, the NSF issued a solicitation requesting responses to and bids on four separate components that represented major changes in the architecture of the nascent Internet. These changes were necessitated by the tremendous growth of the NSFNet over the prior few years, and the push, by private and commercial interests, to allow traffic other than that related to research or education (which was stipulated by the Acceptable Use Policy of the NSFNet). Instead of a single Internet backbone, there would be many, and they would interconnect to each other at Network Access Points (NAPS).

The Internet, as a network of Networks, depends on interconnections as much as it depends on something as basic as network capacity.[1] Every Internet Service Provider (ISP) must inter-operate with other ISPs in order to deliver comprehensive, end-to-end Internet services. No ISP can operate in isolation from others or users on that ISP's network would not be able to reach users on other networks. When the changes actually occurred in early 1995, many issues of these new network to network interconnections, such as which networks should interconnect to others, who should pay and for what, etc., were left unanswered and unresolved. The end result of this indecision was a bill and keep, sender keep all model, in which the network of packet origination charged their customer and kept all of the money, rather than "settle" with the network of packet destination. In the early days, the Internet operators favored network connectivity over all other issues, including short-term profits. Over the past several years, the industry has evolved to the point that most network providers are now trying to attain profitability. The interconnection model

[1] Norton, William. "Interconnection Strategies for ISPs: The Application of a Neutral Internet Business Exchange."

has evolved from an open, non-discriminatory model in which networks connected to each other under nearly any circumstance just to maintain end-to-end reachability, to a hierarchical and somewhat discriminatory nature, in which large ISPs no longer openly peer with all others, but instead, under non-disclosure agreements, try to force smaller ISPs to pay them for interconnections.

If this were the only change, perhaps industry led self-regulation (there have been several such proposals) would be adequate to maintain end-to-end reachability for all users. However, several other issues cloud the picture and make that scenario unlikely. First, the Internet is evolving from a network used for data transfer only to one that is beginning to handle more traditional telecommunications services such as telephony. As a data network, the Internet is classified as an "enhanced service" and is therefore not subject to regulation based on the Communications Act of 1934 and the subsequent updated Telecommunications Act of 1996. Telephone companies and cable companies are beginning to offer Internet services, further clouding the line between the separate services and separate networks. The merging of the traditional, circuit-switched network used to handle telephone calls, the cable television network, and the packet-switched Internet seems very likely.

In light of such "convergence," the Federal Communications Commission may need to re-evaluate its dichotomy between "enhanced services," such as the traditional Internet and cable television networks which are not subject to regulation, and "basic services," which include standard telephone service and is subject to regulation. These issues and their relation to traditionally regulated telecommunications areas such as telephone company interconnection requirements, access charges, reciprocal compensation, the universal service fund, along with other

Internet facets that are facing government scrutiny, may indicate that the future of Internet interconnections is a future of government regulation.

This thesis covers these issues in depth from a United States of America-centric viewpoint. Although Internet interconnection agreements are important all over the world, the state of government regulation in telecommunications is different in nearly every country, so covering them all is obviously not possible. As the U.S. has set an example of Internet development and telecommunications deregulation, focusing on the U.S. should give adequate insight into what may transpire around the world. . In Section I, the origins of the Internet are briefly discussed, followed by an in-depth description of the history of the NSFNet and the issues that led to its dissolution in 1995. Section I also explores the developments of Internet interconnections from the NSFNet dissolution to today. Section II discusses the history of traditional telephony interconnection requirements and settlement mechanisms, and why such mechanisms are much more difficult to define and implement on the Internet. Section III covers the history of FCC regulations over Internet services, as well as current regulatory issues the Internet has created for the FCC, including how the Internet has affected reciprocal compensation and access charge regulations, open access/unbundling requirements in the last mile, the effects of convergence on the universal service fund, etc. Finally, Section IV explores the possibilities of the future of Internet interconnections, by first looking at a few Industry proposed "self-regulated" models and then the possibility of future government regulation. This thesis concludes that the Internet will not be able to live under the umbrella of "non-regulation" due to its current status as an "enhanced service" as defined by the FCC. Instead, the dichotomy between basic and enhanced services will need to be re-evaluated, and in so doing, Internet interconnections will likely be regulated to ensure universal service on the "converged network."

Section I – The History of Internet Interconnections

Section I will explore in detail the history of Internet interconnections. It is important to understand this history to really grasp all of the issues involved in interconnections as well as to see how the industry has evolved to its current state. The origins of the Internet are discussed, from the beginning with the start of ARPANet, to the growth of the NSFNet and its eventual dissolution, up to the Internet as we know it today.

Origins of the Internet

The Internet first started to take shape in the mid 1960's when the United State's Air Force commissioned Paul Baran of RAND to study how it could maintain control over its disperse geographic sites in the aftermath of a nuclear war.[2] In his series of papers "*On Distributed Communications Networks*[3]," Baran proposed a packet switched network with no central hub or control processor, in which packets would be forwarded from place to place until they reached their final destination. The main premise behind this theory was that if any given node were lost due to nuclear bombing (or any other disaster for that matter) traffic would still be able to route around that node and reach its final destination.[4] At that point in time, no such networking hardware or software existed to do anything like what Baran was proposing.

---

[2] Zakon, Robert Hobbes. Hobbes' Internet Timeline v4.0. Available at URL: http://info.isoc.org/guest/zakon/Internet/History/HIT.html

[3] Baran, Paul. "On Distributed Communications Networks." RAND memorandum RM-3420-PR. August, 1964. Available at URL: http://www.rand.org/publications/RM/RM3420/.

[4] Tappendorf, Sean. The ARPANET and Beyond. URL Available at: http://clavin.music.uiuc.edu/sean/internet_history.html. July 30, 1995.

ARPANet

In the late 1960's, the United States Department of Defense commissioned the Advanced Research Projects Agency (ARPA) to further research computer networking,[5] and in 1969, the experimental network called ARPANET was turned online using 56 kilobits per second (Kbps) lines. It was experimental in the sense that it was testing packet switching as a method for data transport,[6] something that had never been done before. In the early to mid-1970's, several issues arose with the network protocol then in use, including problems with the addressing scheme. ARPA assigned the Stanford Research Institute (SRI) to solve the problems with a new protocol. In 1973, Vint Cerf and his group at SRI developed Transmission Control Protocol/Internet Protocol (TCP/IP). In 1976, the ARPANet began experimenting with this new technology which would become the Internet's packet switching protocol and is still in use today. However, it was not until January of 1983 that all computers on the network were required to use it.[7]

During the later 1970's, the ARPANet grew to support many of the organizations of the U.S. Department of Defense and other government agencies, including the National Aeronautical and Space Administration (NASA), the Federal Reserve Board, and the National Science Foundation (NSF).[8] It also began to support both Universities and research organizations. Thousands of hosts, generally those at Universities and government agencies, connected to this network, which had an Acceptable Use Policy (AUP) that prohibited any communications with commercial purposes.[9]

---

[5] Zakon, Robert Hobbes. Hobbes' Internet Timeline v4.0. URL: http://info.isoc.org/guest/zakon/Internet/History/HIT.html

[6] Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html.

[7] Tappendorf, Sean. The ARPANET and Beyond. July 30, 1995. URL: http://clavin.music.uiuc.edu/sean/internet_history.html.

[8] Ibid.

[9] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 3)

The ARPANet continued to exist through the mid 1980's, and in 1984, the Department of Defense separated the military portion of the network, termed MILNET, from the rest of ARPANet. This left non-military government agencies, universities, and research organizations as the organizations that remained on the ARPANet.[10] It was at this point that the National Science Foundation entered computer-networking arena.

NSFNet

In 1984, the National Science Foundation (NSF) created five super computer centers so that people could purchase time on them for research purposes. These five super computer centers were attached to each other at 56 Kbps via a network the NSF established, and the NSF backbone was connected to ARPANET.[11] The five super computer centers were the Cornell Theory Center at Cornell University, the National Center for Supercomputing Applications at the University of Illinois, the Pittsburgh Supercomputing Center, the San Diego Supercomputer Center at the University of California, San Diego, and the Jon Von Neumann Center at Princeton University. Later, the National Center for Atmospheric Research was connected.

By 1985, the ARPANET had become heavily used and congested, and the NSF initiated phase I development of the NSFNet.[12] This network was composed of multiple regional networks and university computer centers Merit, NCAR, BARRNet, MIDnet, Westnet, NorthWestNet, and SESQUINet.[13] Additionally, the two networks NYSERNet and JVNCnet were also connected as each was collocated at a supercomputer center. These components were connected to the NSFNet

---

[10] ARPANET--U.S. DOD-SPONSORED NETWORK, URL Available at: http://www.chaos.com/learn/History.html

[11] Tappendorf, Sean. The ARPANET and Beyond. July 30, 1995. URL: http://clavin.music.uiuc.edu/sean/internet_history.html. July 30, 1995.

[12] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 4)

backbone in a three-tier hierarchy. Campus and research organizations connected to regional networks, which in turn connected to the NSFNet backbone.[14] The original links interconnecting the backbone nodes were 56kbps, and the original routers were LSI-II Fuzzballs.[15]

The NSFNet grew so much that in the fall of 1987 the NSF issued a competitive solicitation to provide a new, faster backbone service to link the six super computer centers and seven mid-level networks.[16] In November of 1987, the NSF awarded a contract to Merit Network, Inc., in partnership with IBM, MCI, and the state of Michigan, to upgrade and operate the NSFNet backbone using T1 (1.544 Mbps) links. This upgrade was completed on July 1st, 1988 (just 8 months after the solicitation), and on July 24th, 1988, the old 56 Kbps NSFNet was shut off.

In 1989, the backbone was reengineered by increasing the T1's so that each site had redundant connections to the backbone, and the router capabilities were increased to handle full T1's. By this time, 500 million packets per month were being switched, which was an increase of 500% in just one year.[17] In January of 1989, the three companies running the NSFNet (Merit, IBM, and MCI) presented a plan to upgrade the network to T3 (approximately 45Mbps) speeds in order to handle the rapidly increasing traffic. IBM developed the first router capable of handling T3 speeds on an RS/6000 subset of Unix. This router would eventually be able to handle 100,000 packets per second. For this upgrade, the NSF wanted a number of new nodes added, and therefore issued a

---

[13] Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html. (p. 1)

[14] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 5)

[15] Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html. (p. 1)

[16] Ibid. (p. 1)

[17] Ibid. (p. 2)

solicitation to anyone who was interested in becoming a new NSF site. The new sites granted were Cambridge, MA's NEARNET, Chicago's Argonne National Lab, and Atlanta GA's SURANet.[18]

In September of 1990, MCI, IBM, and Merit created the company Advanced Network and Services (ANS).[19] Merit provided the policy routing database, routing consultation, and management services of NSFNet, whereas ANS operated the backbone routers and had a Network Operations Center (NOC). Basically, ANS was created to operate the 45Mbps backbone. IBM and MCI each gave $4 million to the network and ANS acted as a subcontractor to Merit. The final T3 was installed in November of 1991, and at that time the network had 16 sites with total of 3500 networks connected.[20]

The Beginning of Commercialization and Privatization of the NSFNet

In early 1992, just after the upgrade to T3 had been completed, a contentious debate about the evolution and commercialization of the US Internet began.[21] ISP's, including those that provided local dial-up services and those that provided T1's and TCP/IP products and services, were springing up all over the country. In addition, the NSFNet backbone had reached "critical mass" – enough so that the problem of "email islands" could be over come.[22] Before the NSF had reached this level, users in the commercial domain that wanted to send email were limited to sending to those users attached to the same network. With the ability to send email from network to network,

---

[18] Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html. (p. 2)

[19] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 4)

[20] Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html. (p.3)

[21] Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html. (p. 2)

the NSFNet grew even more. There was such a large user population, and participation was so high, that it became a thing in itself to connect your private network to the NSFNet. And the more small networks that attached to it, the more attractive it became to get attached. At this early date, THE Internet was defined as having connectivity to the NSFNet. [23]

Between 1988 and mid-1990, an Acceptable Use Policy (AUP) had required that all backbone traffic "further research and other scholarly activities."[24] The language of the AUP changed slightly in 1990 such that the NSFNet's purpose was redefined "to support research and education in and among academic institutions in the U.S. by [providing] access to unique resources and the opportunity for collaborative work.[25]" However, by 1992 there were many pressures being felt on the NSFNet for commercial and general-purpose groups that wanted network access[26]. ISPs were emerging to accommodate this need, thereby defining a new industry, while international networks had developed and also wanted to connect. The Internet was no longer an experiment in computer networking, and much of the Internet's traffic had little to do with anything the National Science Foundation was charted to do. The obvious solution was to privatize it – to allow commercial companies to operate it and sell access to it.[27]

Discussions relating to the commercialization and privatization of the Internet began as early as 1990. The NSF and the U.S. Congress's Office of Technology sponsored a workshop from March 1st through 3rd of 1990 entitled "Commercialization and Privatization" which was held at the

---

[22] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes. May 1998.

[23] Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html. (p.3)

[24] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

[25] Ibid.

[26] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 4)

[27] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes, May 1998

John F. Kennedy School of Government, Harvard University. This workshop explored the issues of commercialization of the NSFNet and published a report entitled "Commercialization of the Internet Summary Report" as RFC 1192.[28]

Stephen Wolf, then Director of the NSF Division of Networking and Communications Research and Infrastructure, distinguished "commercialization" and "privatization" on the basis of his experience in deploying policy for the NSFNet as follows:

> **Commercialization:** permitting commercial users and providers to access and use Internet facilities and services
>
> **Privatization:** elimination of the federal role in providing or subsidizing network services.

The resulting publication of the conference (RFC 1192) suggested that "offering the NSFNet backbone at no cost to authorized networks both encourages undisciplined use of the backbone and inhibits private investment in backbone networks." The workshop provided several options of what to do, including charging the regional networks or moving the subsidy down to the campus level or the researchers themselves. It came down to funding the backbone or funding the researchers, institutions, or regionals. Privatization might be achieved by shifting the federal subsidy from network providers to users, thus spurring private sector investment in network services, but there were several issues with this idea. There were concerns over the acceptable use policy, but with a private/commercial network, the nature of the traffic would not have to be bound to the NSF's charter. Also, at that time, there were several commercial services such as Lexis, Dialog, and Orbit, but they were not widely used and the services themselves were generally unaware of their own potential. The conference concluded that "without a critical mass of users, commercialization may

[28] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

need to precede privatization." There were no real actions drawn from the conference, but it should be remembered that the conference was held in 1990.

Back to 1992: there continued a great discussion regarding the commercialization and privatization of the NSFNet.[29] Private networks paid for by private monies were starting to grow, and many sought direct or indirect access to the NSFNet. (The AUP did not restrict such connections, if, in one-way or another, the added network could support the mission.)[30] Should the government fund and operate something that competed with private companies such as Sprint, MCI, UUNET, CompuServe, Performance Systems International (PSI), and others, when a growing amount of traffic was becoming more and more commercial? In some senses, the network was already private – or at least certain aspects of it were. The physical circuits were owned by the private sector, and the logical network was managed and operated by the private sector. Nonprofit regional networks attached to the NSFNet were increasingly contracting out routine operations such as network operations centers (while retaining control of policy and planning functions).[31]

At the backbone level, the NSF was contributing less than $3 million of an annual $10 million in expenses. The State of Michigan gave $1 million, and the rest was covered by the subcontractors Merit, IBM, and MCI. At the regional level, the NSF was covering 40% of operating costs. There was no federal funding of the Campus level networks, which were estimated to be seven to ten times larger in annual investments than the mid-level networks.[32]

The NSF allowed the NSFNet to connect commercial users and relax the strict limits of use on the backbone. In 1992, Representative Rick Boucher's subcommittee on Science, Research, and

---

[29] Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html. (p.3)

[30] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

[31] Ibid.

[32] Ibid.

Technology asked the Office of the Inspector General to conduct an investigation into the legalities of the NSF's move towards privatization. The Office of the Inspector General concluded that the NSF did in fact have the power to relax the AUP, and Rep. Boucher then attached an amendment to the NSF Authorization Act of 1993 to explicitly state so.[33]

With the push towards commercialization and privatization, there were many different views on what should happen. "IBM and MCI had essentially built the NSF backbone, and claimed to have an identical private backbone, using the same rooms, the same equipment, and the same technicians as the NSF backbone."[34] Both companies thought they should inherit the backbone because of their considerable investments in it. However, there was a small group of outspoken critics, including Rick Adams of AlterNet (now UUNET, one of the largest Internet Service Providers and now owned by MCI Worldcom), Marty Shafkowits and Bill Schrader of Performance Systems International (PSI, still a large ISP) and several others. Their view was that a government program funded by tax dollars should not be given to a private company, and they made sufficient noise such that they too were allowed to participate in the coming privatization.

Commercial Internet eXchange (CIX)

Before the debate began, back in 1991, private network operators including PSI Net, General Atomics (renamed CerfNet, then bought by TCG, and then by AT&T), and UUNET formed an exchange point called the Commercial Internet Exchange (CIX) in UUNET's Falls Church, VA offices to help facilitate commercial networking. CIX was created due to a mixture of events.

[33] Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Forthcoming from the Cato Institute under Telecommunications Policy Analysis: http://www.cato.org/research/telecom-st.html.

[34] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes, May 1998

ANS had threatened to charge networks for connections to the NSF backbone[35] and the increasing burden commercial traffic was having on the NSFNet (despite the AUP restricting such traffic) had some networks concerned. The goal of CIX was not necessarily to relieve the traffic loads on the backbone, but to instead facilitate the exchange of traffic between private networks. In fact, most commercial traffic still crossed the NSFNet backbone even after CIX had been created.[36] However, the CIX was an early start on commercialization.

CIX had many important implications for the future of the Internet as described here. First, there was a fundamental agreement that all CIX members would interconnect with one another and there would be no restrictions on the type of traffic that could be routed between each network. Also, there would be no "settlements" or any traffic-based charges between CIX member networks. Basically, this meant no network could charge another for interconnecting or the traffic that was exchanged.

One obvious benefit for CIX members was that as the networks connected to each other, they instantly had access to each other's users and services, greatly increasing the value of their network connection. Of course, this was true of connecting to the NSFNet as well; except with the NSFNet, there were AUP restrictions on what the nature of the traffic was allowed to be. (It should also be noted that CIX had other benefits to members that are not directly related to the evolution of Internet interconnections. One such example is that it provided a neutral forum to develop consensus positions on legislative and policy issues.)

Perhaps the most important aspect of CIX, and one that was most likely not noticed at the time, was that it was the 1st point of interconnection for the networks attached to the NSFNet

---

35 Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Forthcoming from the Cato Institute under Telecommunications Policy Analysis: http://www.cato.org/research/telecom-st.html.

36 Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html.

backbone that did not actually cross the NSFNet backbone. This would become an early model for Internet Interconnection as the NSFNet was dissolved.

Other Developments

Similar to the CIX, except for Federal institutions such as the NASA Science Internet (NSI), the Department of Energy's ESNet, DARPANet, etc., the Federal Internet eXchange was created to exchange traffic with member networks at a place other than the NSFNet backbone.[37] In addition, on the private sector side, MFS DATANET established the Metropolitan Area Ethernets (MAES) in Washington, D.C., and San Francisco. These two developments are important because of the precedence they set as exchange architectures. Rather than each network attaching to the NSFNet backbone in order to exchange traffic with one another, they could interconnect directly with one another, and other networks, at a shared facility.

Dissolution of NSFNet

"The early years of the Internet, under the management of the NSF, were characterized by sharing, bound together in part by technical necessity."[38] With a single common backbone and lack of private interests trying to be profitable, there was a common bond to create a network that everyone could easily attach to. However, the network grew, and a new architecture had to be

[37] NSF 93-52 – Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNET and the NREN (SM) Program. May 6, 1993. National Science Foundation.

[38] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998. (p. 3)

created. The sharing and openness of the early days would continue through the dissolution of the NSFNet, but only for a short time beyond that.

Why the solicitation was needed

The NSFNet had started out serving its original purpose well – it had established a backbone network that regional networks connected to, which in turn attached to universities and government networks. They all followed the NSF's charter and Acceptable Use Policy (AUP), which was to "further research and other scholarly activities" and later "to support research and education in and among academic institutions in the U.S. by [providing] access to unique resources and the opportunity for collaborative work.[39]" The NSF had supported the data networking needs of research and education since 1986 and had become an essential infrastructure for that community. The number of users and the number of connected networks, along with the amount of network traffic continue to grow rapidly. By 1993, the network of networks, which had just started to be called the Internet, connected 10,000 IP networks, which had more than one million computers and millions of users throughout the world.[40] Data networking had evolved considerably since the inception of the NSFNet, and new companies had been created while a number of existing companies had shown increasing interest in the field.

For these reasons, and due to the expiration of the Cooperative Agreement for the NSFNet Backbone, a new architecture was needed. The NSF formulated specifications for such a design in a solicitation entitled *"NSF 93-52 – Network Access Point Manager, Routing Arbiter, Regional*

---

[39] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

[40] NSF 93-52 – Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNET and the NREN (SM) Program. May 6, 1993. National Science Foundation

Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNet and the NREN (SM) Program."[41]

The Solicitation

NSF 93-52, issued in May of 1993, would radically alter the architecture of the Internet. In short, the NSF was getting out of the backbone business. The solicitation had four main points as described here:

1. Create a set of Network Access Points (NAPs) where major network providers would connect their networks, exchange route advertisements, and exchange traffic.

2. Implement a Route Arbiter (RA) project to facilitate the exchange of policies and addressing of multiple providers connected to the NAPs.

3. Find a provider of a very high-speed Backbone Network Service (vBNS) for educational and governmental purposes.

4. Transition existing and/or realigned regional networks to support interregional connectivity by connection to NSPs that would be connected to NAPs or by connecting directly to NAPs. Any NSP selected for this purpose had to be connected to at least 3 of the NAPs. (An NSP was defined as "a network service provider that connected to all of the NSF priority NAPs and provides, at a minimum, IP internetworking service at DS1 speeds.")[42]

The focus of this paper is most concerned with number one, the creation of NAPs where routing updates and traffic would be exchanged. This was the true start of "Internet

[41] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

[42] Program Solicitation NSF 92-52, Questions and Answers. June 23, 1993. National Science Foundation. (p. 3)

interconnections," where different networks connected to each another directly, rather than the prior model where the NSFNet backbone was the "intermediary" connection point. (See Diagrams 1 and 2) Part four of the solicitation is of secondary concern in this paper, as the regionals would have to interconnect to NSPs in order to remain "connected" to the Internet.

Diagram 1: How ISP's connected to each other before NAPs were created. Here, the NSFNET was a nationwide network run by the NSF:

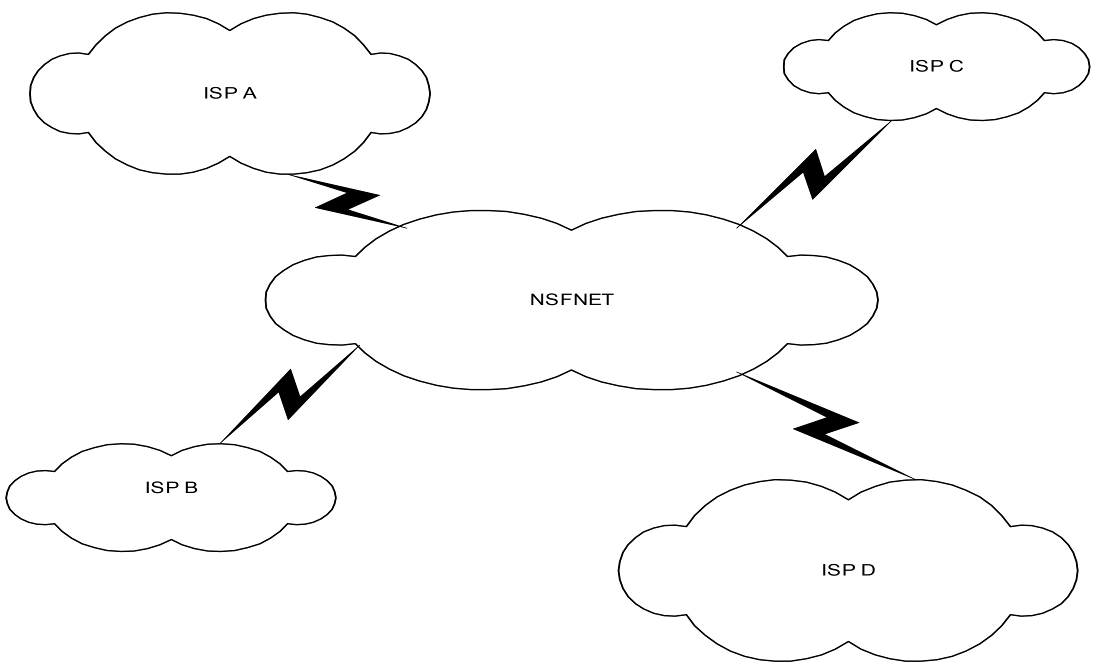

Diagram 2: How networks attached to each other after NAPs were created. ISPs would generally connect at two or more geographically disperse NAPs (not shown here). The ISPs could be small regional sized networks are large national networks.

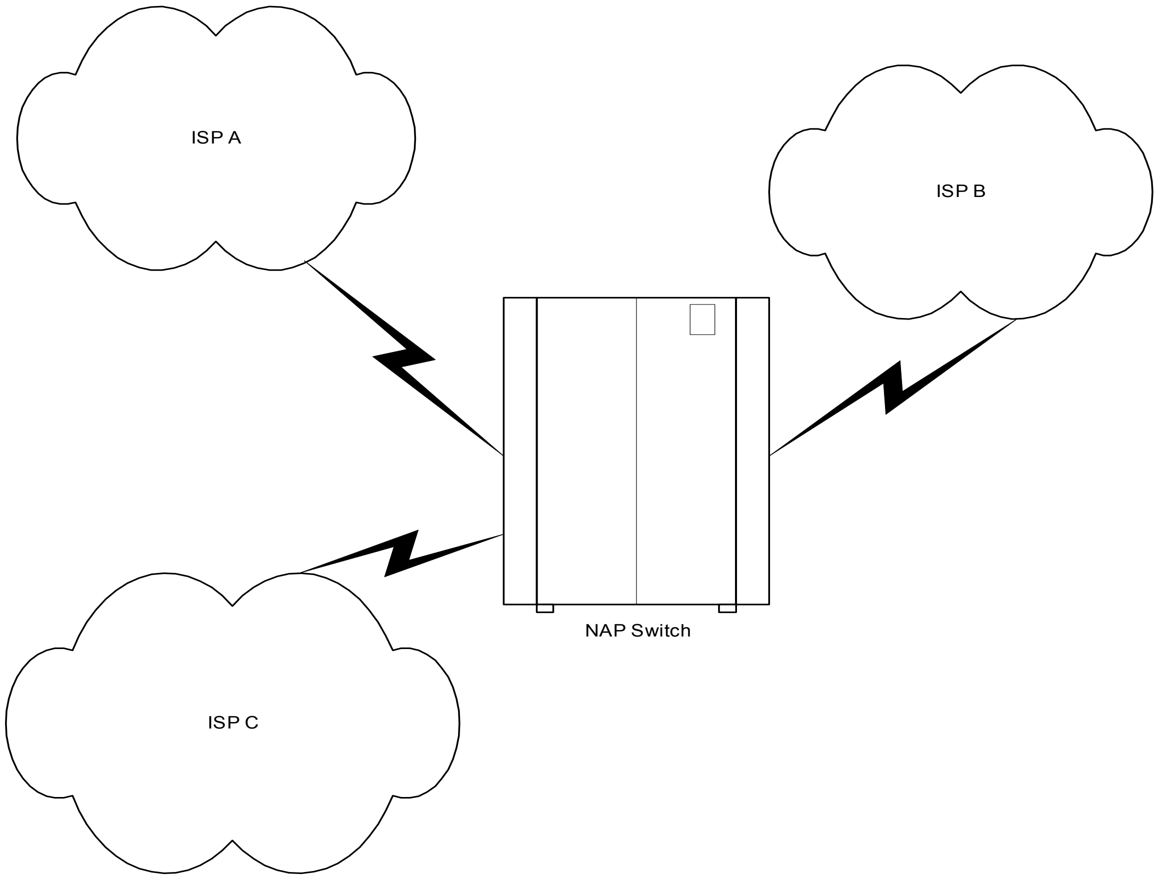

As previously stated, the solicitation sought to totally transform the nascent Internet. As it stood, the NSF backbone provided the Interconnection for all of the networks, except for the CIX and FIX interconnection points. However, after the implementation of the solicitation, there would no longer be a single backbone. Instead, there would be several national backbone providers (the NSPs) that would interconnect at the NAPs, which were quite similar in concept to the CIX and FIX. The move was away from a core backbone to a distributed architecture operated by commercial providers.[43] The due date for responses to the solicitation was August 17, 1993.

[43] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 4)

Network Access Points (NAPs)

NSF 93-52 defined an Internet NAP as "a high speed network or switch to which a number of networks can be connected via routers for the purpose of traffic exchange and interoperation." The solicitation further specified that NAPs should "have adequate capacity to keep up with the switching requirements of the attached networks" and that the "attached networks are presumed to be a part of the connected Internet, but the NAP itself may be of a lower protocol layer; e.g. it may be a level two network or switch." The solicitation went so far as to say that the NAP would be a conceptual evolution of the FIX and CIX Internet exchange points, and gave examples of what the NAP switching fabric could be: LAN, MAN (using SMDS), or a high speed switch such as ATM.

Another important aspect of the solicitation for NAPs was that traffic on these new exchange points would not be restricted to the NSFNet backbone's AUP (e.g. "that which is in support of research and education"). In other words, there would no longer be restrictions on the type of traffic that could cross the network. Therefore, this was the true start of commercialization and privatization of the Internet.

NSF 93-52 specified three priority locations for NAPs – California, Chicago, and New York City. It further listed desirable locations as Atlanta, Denver, Boston, Texas, and Washington, D.C.

The solicitation further defined what the anticipated duties of the NAP manager organizations would be, as listed below:

- Establish, operate, and maintain … all or a subset of the specified NAPS.
- Develop and establish attachment policies, including attachment fee schedules, which should be fair and equitable. The fees were to be proposed by NAP managers in their response to the solicitation, but final approval would be granted by the NSF.
- Propose NAP locations subject to the priority and desirable cities listed in the

solicitation.

- Specify reliability and security standards for the NAPs.
- Specify and provide appropriate NAP accounting and statistics.

In the end, the NSF awarded four NAPs as follows:[44]

1. Sprint NAP – Pennsauken, NJ
2. PacBell NAP – San Francisco, CA
3. Ameritech Advanced Data Services (AADS) NAP – Chicago, IL
4. MFS DataNet (MAE-E) NAP – Washington, D.C.

The NSFNet physically attached to Sprint NAP in September of 1994; the PacBell NAP in October of 1994, to the AADS NAP in January of 1995, and to MAE-E in March of 1995. MAE-E had actually been doing Internet interconnections on a significant basis before the other three NAPs were up and running and the NFS grant to this NAP was contentious[45]. MAE-E was only a 10Mbps Ethernet, and the solicitation had listed a requirement of 100M or more. MFS did upgrade to 100M Fiber Distributed Data Interchange (FDDI), and, though there was still some contention due to the fact that ATM was the solicitation's desired switching fabric, MAE-E was never-the-less granted the status of an official NAP. The fact that MAE-E was run by MFS, a neutral 3rd party, was also seen as an attraction.[46] The other 3 NAP managers were not quite as neutral as MFS was. PacBell and Ameritech were local telephone companies that at the time were just entering the Internet access business, whereas Sprint was a major provider of Internet backbone services.

---

[44] Halabi, Bassam. 1997. Internet Routing Architectures. Cisco Press. (p. 12)

[45] Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Not yet published.

[46] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

Importance of NAPs

NAPs provided the foundation of network interconnections that allowed the Internet to be just that – a network of networks. When they were first created, they provided a logical, scalable, and cost effective way to connect networks. Private backbone networks could interconnect and exchange traffic directly at the NAPs, rather than indirectly on an intermediary backbone such as the NSFNet.

The geographic location of NAPs was (and still is) quite important. First, the more interconnections points between two networks, the more efficient traffic could flow. For example, in diagram 3 below, ISP A and ISP B only interconnect at one point, so all of there traffic must pass through that point. This is both inefficient and can lead to major outages if there is a problem at that one point.

Diagram 3:

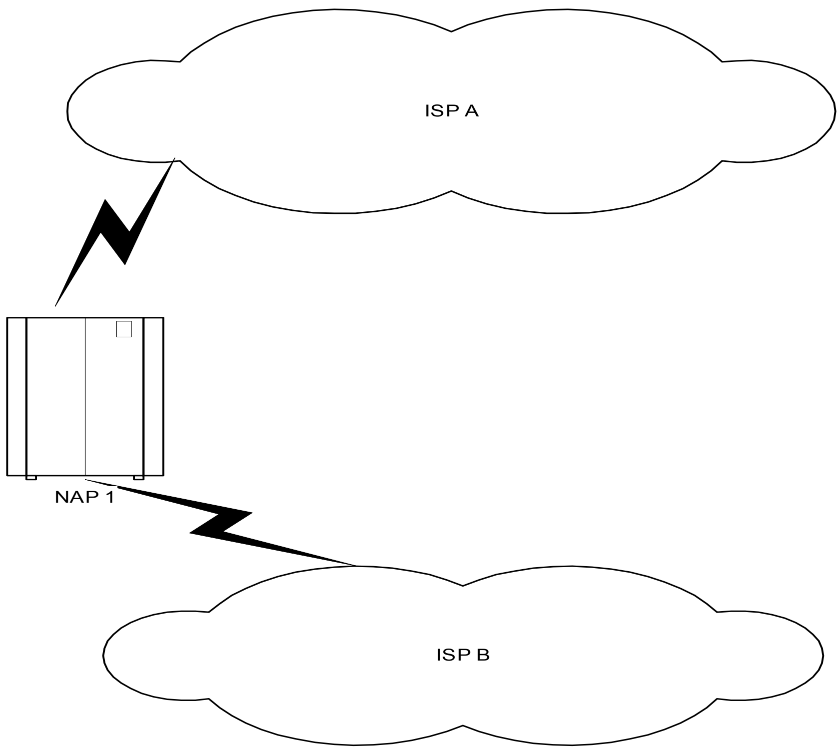

If ISP A and ISP B interconnected at two points, as shown in Diagram 4, traffic flow could be more efficient and there would be more redundancy.

Diagram 4:

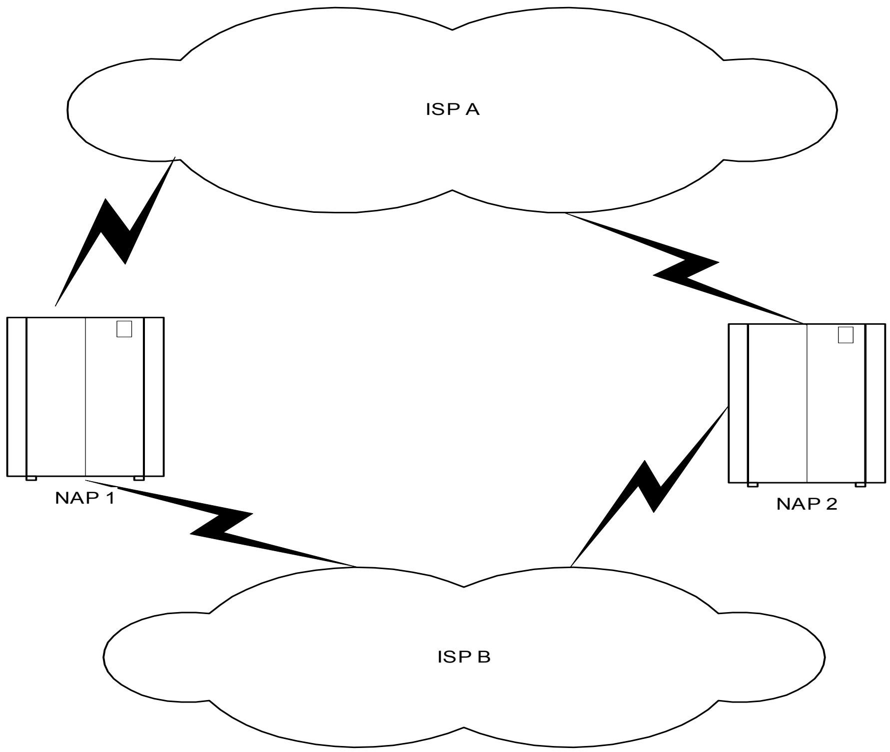

The Internet evolved around shortest exit routing (sometimes called hot potato routing). What this means is that traffic originating on network A should be sent to the destination network B at the "nearest" interconnection between A and B. Therefore, the more NAPs there are near the origination and/or destination, the less backbone infrastructure for each single network is needed. This is depicted in Diagram 5, in which the NAPs have been replaced with simple communications links for simplicity. In this case, packets from the end user to the destination computer would go over interconnection point 1, since that is the "closer" link. Packets from the destination computer back to the end user, would travel over interconnection point 2, because in that direction, that link is the closest interconnection point between the two networks.

Diagram 5:

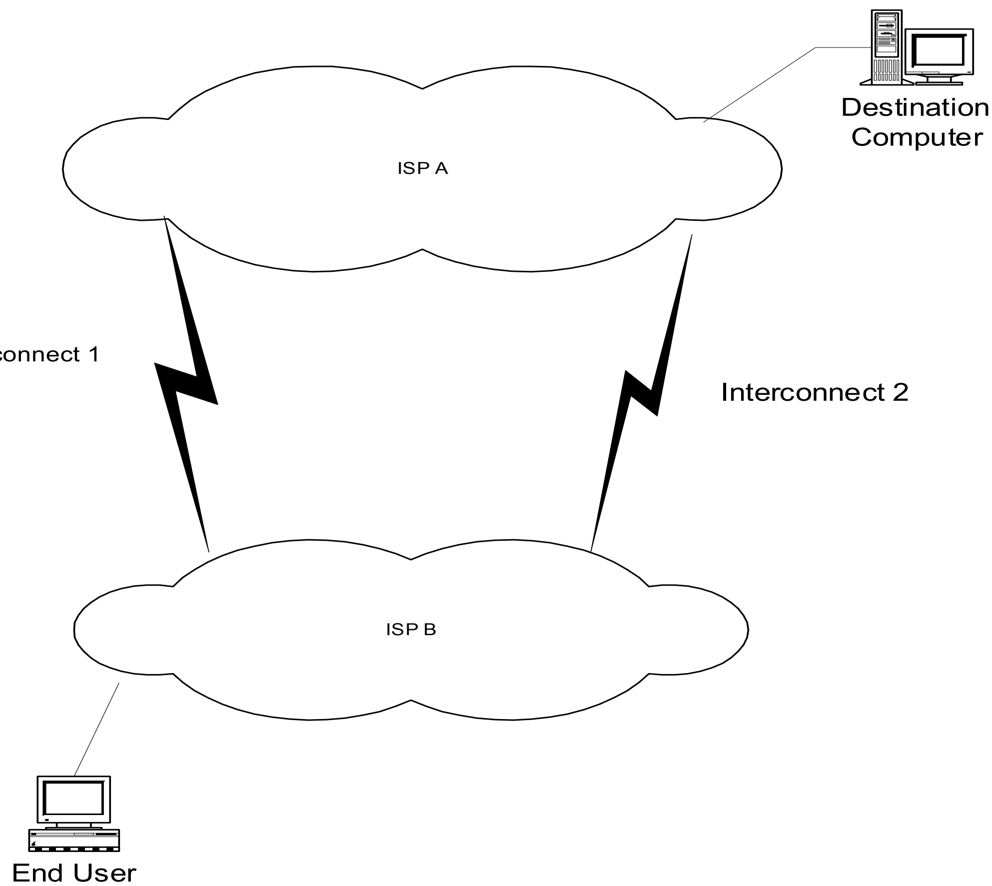

Shortest exit routing and the existence of several NAPs also established the fact that interconnections were good thing. A good way to understand this is with a counter example. If network A and network B only interconnect at the MAE-E NAP, then a network A end-user in California that wants to reach a network B server that is also in California, would actually have to pass traffic all the way to the east coast at the MAE-E NAP in Washington DC. If networks A and

B also interconnected at the PacBell NAP in San Francisco, then the traffic between the two end users in California would remain relatively local. Note that this fact means the level of redundancy and diversity of NAP connections affects traffic patterns of whole Internet.

It should be noted that there is a limit to “the more interconnections between two networks the better” philosophy. There are issues with scalability and the complications many interconnections can add to network configurations, management, operations, etc. There are also economic scalability issues as each interconnection costs money to install, maintain, and operate.

Although NAPs were free from the NSF’s AUP that restricted what kind of traffic could be passed, they did have one important restriction. The traffic passed at the NAP had to be “backbone to backbone” traffic only – i.e. no “transit” traffic could be passed. Transit traffic can be difficult to clearly define. Generally, a given service provider “buys transit” from a national service provider, to pass all their traffic to the rest of the Internet. There are anywhere from 10-15 “transit-free” providers, meaning that they pay no other network to have “global” connectivity. This paying relationship is what was prohibited at the NAPs, such that network A could not pay network B to carry their traffic when the interconnection between A and B was handled over the NAP infrastructure. Network A could buy transit from Network B, but their interconnection had to be at private location. This is depicted in diagram 6.

Diagram 6:

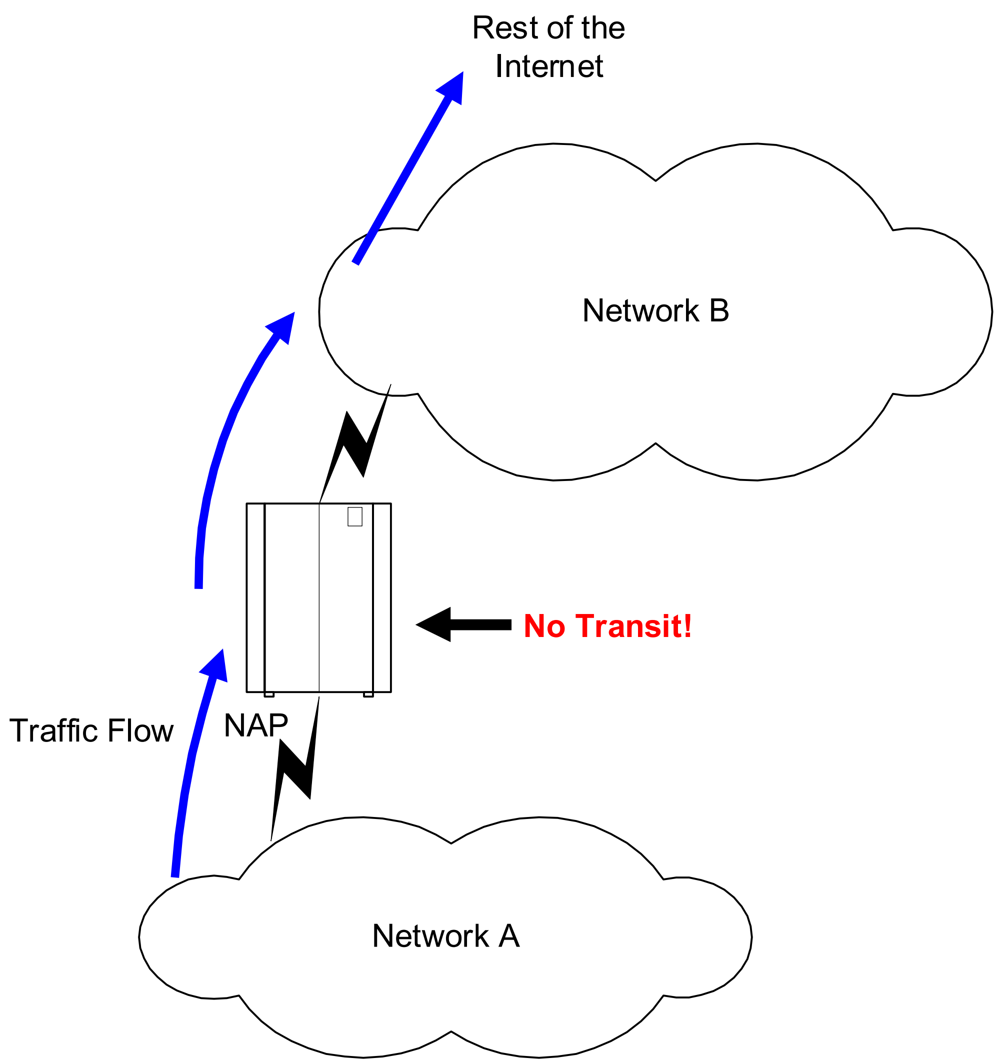

NAPs continue to be the model of interconnection around the world as the Internet pushes its international boundaries. As countries with more than two Internet Service Providers connect to the Internet, a NAP architecture is generally the most cost effective, simplest way for those providers to interconnect within their own country, rather than long hauling traffic back to the US.

Regionals

It is important to briefly consider the regional providers at this point, as their future in Internet interconnections will be relevant in later discussions. The regional networks were a part of the NSF since its inception and were a major force in providing ubiquitous network connectivity for the research and education communities.[47] The NSF 93-52 solicitation recognized this role and explained in detail what should happen. The NSF expected that most regionals would elect to use an NSP to obtain inter-regional connectivity and connectivity to the NAPs.[48] The regionals could request funding for both NAP and NSP connections from the NSF (with some restrictions), though NSF financial support for such connections was to decline over 4 years.

The regionals that had to be moved to an NSP or NAP included all of those that had been directly connected to the NSFNet backbone, such as SURANet and NYSERNet, but also included other down stream networks such as NevadaNet and MoreNet. Most regionals selected internetMCI or SprintLink as their NSPs to connect them to the rest of the Internet.[49]

The deadline for moving their connections from the NSFNet to an NSP was the end of October 1994, but the task proved to be more difficult than first thought. In mid-October 1994, NSFNet Program Director Priscilla Huston sent a letter to the regionals asking them to send a transition calendar and engineering overview to Merit and also asked the regionals to notify her if they weren't going to make the October 31st deadline for moving off NSFNet."[50] No regionals were able to make that date, as most pushed back due to provisioning delays the NSPs had. In one or two cases, the regional had to back out of a move to an NSP back to the NSFNet due to connectivity problems. By mid April of 1995, only seven regionals had completely disconnected

---

[47] NSF 93-52 – Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNET and the NREN (SM) Program. May 6, 1993. National Science Foundation. (p.2)

[48] Program Solicitation NSF 92-52, Questions and Answers. June 23, 1993. National Science Foundation. (p.3)

[49] Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html. (p. 6)

from the NSFNet. However, by the end of April, all of the regionals had been moved and the NSFNet was officially shut down.[51]

No Interconnections Guidelines Given

One major issue that the NSF did not address during the NSFNet backbone shut down was all of the details of Internet Service Provider interconnections. The NSF defined NAPs as places where ISPs could interconnect, and it suggested how the regionals could connect to the Internet via the NAPs and/or NSPs. However, NAPs only provided a place for interconnecting to other ISPs, where any one network could choose to interconnect with another network. Nevertheless, those NSPs could also choose to not interconnect. It was and still is relatively simple to get a connection to any NAP, but it is an entirely different and much more difficult thing to get other networks to peer with you there. There was much debate on this subject just prior to and just after the dissolution, but no agreements were ever reached.[52] This lack of decision would have profound impacts on the future of Internet interconnections, and is therefore at the heart of this thesis.

From the NSFNet breakup until Today

Although the NSFNet was officially dissolved in April of 1995, and the Internet architecture changed dramatically then, it has continued to evolve over the last several years. This section will

---

[50] Ibid.

[51] Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. URL http://boardwatch.internet.com/isp/internetarch.html. (p. 4)

[52] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes. May 1998. (p. 2)

explore those changes in peering and interconnections and why they have occurred. In order to understand some of the discussion, some service provider definitions are given first.

Service Provider Definitions

At this point, some definitions of common terms should be given so that the following discussion can be clearly understood. Generally, there are three levels of Internet Service Providers, though there are no agreed upon exact definitions and some providers may fit the definition of more than one such category:[53]

- **Tier 1 ISP:** Tier 1 ISPs have extensive backbones capabilities, at least on a national, if not international scale. They are connected to other tier 1 providers at several locations in the U.S. and sometimes in other geographies as well. Tier 1 often implies that the provider is able to get all routing information about the Internet from other tier 1 providers and does not have to pay for transit connections to anyone. Transit is defined below.

- **Tier 2 ISP, Regional/mid-size:** Some ISPs in this category are often national in a sense, but they do not own significant backbone facilities, and therefore depend on providers that have more backbone presence. Other ISPs in this category are truly regional, tending to provide service in one particular location. (That location may be of large size, such as the Northeastern United States, the Mid-West, etc.) Regional providers typically connect at one NAP or with one tier 1 provider. In areas where they may have a backbone presence, they may peer with a tier 1, yet in other areas, they may buy transit services.

- **Tier 3 ISPs, small ISPs, and resellers**: ISPs in this category generally resell another ISPs service. Regional Bell Operating Companies (RBOCs) are sometimes said to fit in this category, as they can not legally transport data over prescribed

[53] The definitions listed are paraphrases from "The Need for Facilities-Based Internet Backbone Competition" by Robert Gibson, though information from many sources was used to craft them.

service boundaries known as LATAs. However, in another sense, RBOCs are regionals because they do have regional presence.

- **Peering/interconnection:** In general, peering and interconnection are used interchangeably throughout this paper. These terms describe a relationship between two service providers in which they exchange their routes and their customers' routes with each other, and neither is paying the other for this service. In other words, there is no customer / provider relationship between the two service providers. Such interconnections can be at public NAPs, or could be over private interconnects which will be discussed more below.

- **Transit:** In a transit relationship, one provider pays another to carry traffic to destinations that they don't get from any other peers. A transit customer can be as simple as a customer that buys a single circuit to a single site. In this case, all of their traffic must transit their provider's backbone to reach the rest of the Internet. A more complicated example might be a regional provider that has peering connections to other regionals in the same geographic region, but buys transit from one or more national providers to get out to the rest of the Internet.

The Congestion of the NAPs and other Issues

The NAP model was all that existed following the NSFNet backbone dissolution. National Service Providers connected to the NAPs, and via the NAPs to each other so that they could exchange routing information and traffic.

The growth of the Internet has been exponential since the early days of the NSFNet backbone. The Graphical, Visualization, and Usability Center at the Georgia Institute of Technology kept NSFNet traffic statistics up until the dissolution of the backbone, and also projected traffic growth from the historical data into the short term future. This information is available at URL:

http://www.cc.gatech.edu/gvu/stats/NSF/merit.html

In addition, the MAE-East NAP, which is generally the largest NAP in terms of traffic, publishes their statistics from January of 1997 until now, at http://www.mae.net/east.stats.html. These graphs also show Internet traffic growing at exponential rates.

The NAP model uses a shared layer 2 switching fabric, and this architecture has struggled to keep up with Internet growth. Layer 2 of the Open Systems Interconnection (OSI) model is the data link layer (R36). Ethernet and FDDI were the early media used by the NAPs, followed by ATM (which covers several layers of the OSI model). What is meant by "shared layer 2 switching fabric" is that each provider connects to a shared device at the NAP. FDDI switches, or Ethernet (now Fast Ethernet) switches, and ATM switches, all allow many providers to link to each other at the data link layer.

In some cases, as the NAPs grew, one switch was not enough, so more had to be added, and then the multiple switches of course had to connect to each other. The growth of traffic of the Internet has far out-paced the ability of this model to scale, and it has caused packet loss and latency for Internet traffic.

> As traffic levels increase, the queues on network interfaces (i.e. on routers and switches) lengthen. As the queues grow longer, the time a packet waits in the queue grows, and latency increases. As network hardware resources are finite, at some point the queues fill, and packets get dropped. [54]

[54] This definition is a paraphrased version of the Cross-Industry Working Team Internet Protocol Performance Measurement working group's definition as given in "Customer View of Internet Service Performance: Measurement, Methodology, and Metrics." Cross-Industry Working Team Internet Protocol Performance Measurements workgroup. Available at URL: http://www.xiwt.org/documents/IPERF-paper.pdf.

When an ISP has packet loss and/or latency to another ISP because the NAP has not been able to scale to the demand of traffic between the ISPs, all users of both ISPs can suffer.[55] The NSF solicitation 93-52 did not explicitly define requirements of when NAP infrastructure must be upgraded, and to be fair to the NAP managers, the growth has been nearly impossible to keep up with as the technology needed to support the required traffic levels has often been "a month too late." These issues eventually led to hybrid-private peering and later private peering, which will be discussed below.

There were other problems with the NAPs that have slowly caused some changes to occur. For example, recall from the "Importance of NAPs" section that the geographic distribution of the NAPs means that packets have to be "long-hauled" in many cases. "Long- hauled" means that even though the two computers communicating may be close to one another, the packets between them may have to travel long distances to the point where the two computers' network providers interconnect. As the four main NAPs were located in San Francisco, Chicago, New Jersey, and Washington DC, traffic that originated, for example in Florida on a computer attached network A and was destined for computer in Florida that was attached to network B another would actually have to travel to the nearest NAP in Washington DC. This is depicted in diagram 7 below.

Diagram 7:

[55] Huston, Geoff. "Interconnection, Peering and Settlements – Part I." Internet Protocol Journal, Volume2, Number 1, March 1999. (p. 14). Soft copy available at URL http://www.cisco.com/ipj

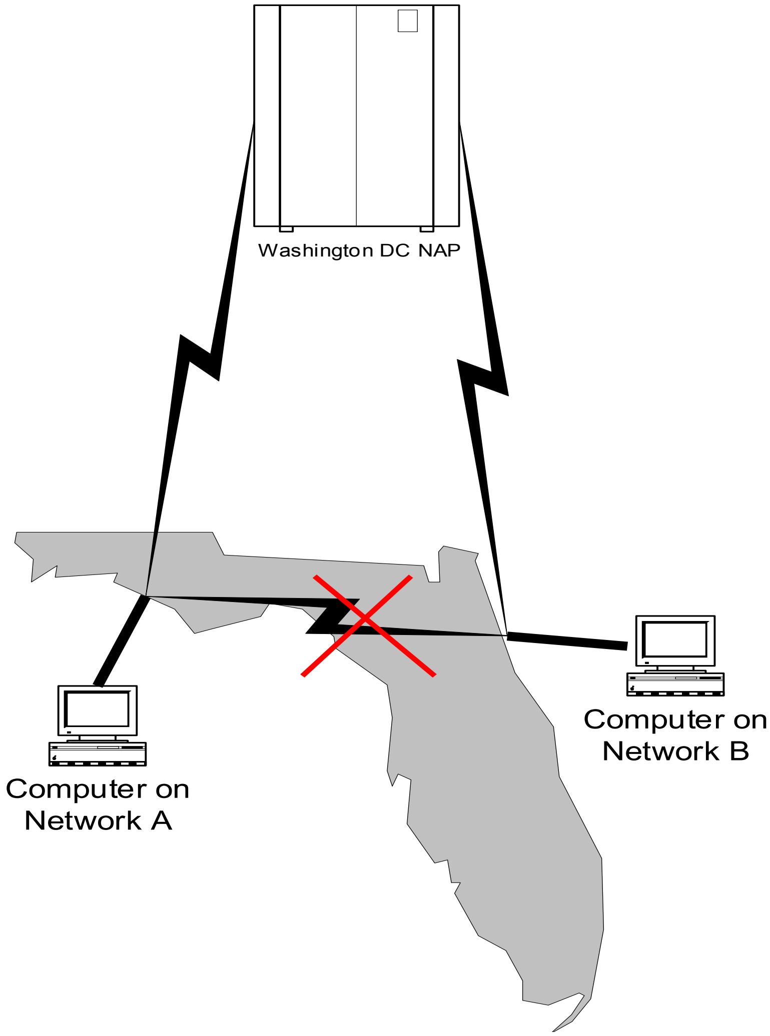

Another issue with the NAPs is that they are often owned by a party with a vested interest in their own ISP. Recall that the neutrality of the MFS operated MAE-E was seen as a benefit at the time the NSF named it as an official NAP. However, MFS was bought by Worldcom, which also owns UUNET, the worlds largest ISP.[56] Similar situations occur at the other NAPs, with AADS in Chicago being owned by Ameritech which runs ISP services, and with the Sprint NAP in New Jersey, owned by Sprint, which runs an ISP service. To connect to a NAP, an ISP is forced to by a circuit to the NAP from (in most cases) a single telephone company, which is often the same

[56] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998. (p. 3)

company that owns the NAP.[57] [58] This makes it difficult to get good pricing on the circuits, as there is no real competition for circuits to the NAPs.

Hybrid Private Peering

Due to the heavy congestion at the public NAPs, which caused latency and packet loss, hybrid private peering began to appear. This model of interconnection bypassed the NAPs' congested switches by using "private" wires between two ISPs' routers at the "public" NAP, and moving all traffic exchanged between those two ISPs to that private wire. This allowed for better connectivity (less loss and less latency) between those two ISPs, while also alleviating the NAP switch problems somewhat, so that those ISPs still using the switches got slightly better performance.

Some NAPs were reluctant to allow this, some charged for the service, and some NAPs, such as the Palo Alto Internet Exchange (PAIX) were actually built around this model. In addition, a new series of NAPs, such as those being built by Equinox, share this model. In certain types of NAPs, no collocated router is necessary. As collocation of a router at the NAP costs more money due to space requirements, there are economic reasons to not have a router on site. However, if a router is not on site, there is no possibility of implementing hybrid private peering.

---

[57] Norton, William. "Interconnection Strategies for ISPs: The Application of a Neutral Internet Business Exchange." And

[58] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998. (p. 31)

Emergence of Private Peering

In some cases, hybrid private peering was either unavailable or it introduced other problems. For example, if a given ISP had several direct cross connects to other ISPs at the NAP, in addition to other connections to other ISPs over the normal NAP infrastructure, traffic levels could increase so much that the leased line from the ISP's backbone to the NAP could become over utilized. As previously discussed, over utilization on any network component can lead to increased latency through that device, and eventually to packet loss as interface queues fill up.

Large ISPs began to interconnect privately over direct circuits (in contrast to shared switches) in places mutually convenient for each other. Such private circuits provide several benefits. First, they are often faster than the congested NAPs. Second, as there is no third party provider of the interconnection fabric, operations (maintenance, repair, troubleshooting, etc.) are simpler. Also, the two providers have more control over the peering agreement, the configurations, and the available bandwidth.

These benefits do not necessarily come without any drawbacks. For one, the private circuits cost more. At a public NAP, an ISP has to buy a single circuit to the NAP, but there are many interconnections over the shared NAP infrastructure. However, with private circuits, there must be a single line for each peer. Most peering agreements are covered by non-disclosure agreements (NDA's) and therefore, many details regarding the legalities of the interconnection, including which party pays for what are not known. However, in general, the cost of direct circuits between two ISPs is shared. ISPs will generally peer in at least two locations for reasons of redundancy and in order to not long-haul traffic, and each of those circuits is paid for by one of the ISPs. This is one arena in which facilities based ISPs (i.e. companies that own fiber optic networks) have a large advantage in private peering. They are able to provision the circuits quickly and operate them at lower costs, whereas an ISP that has to lease the circuit from a telephone company often has to wait long periods of time for the provisioning and has to pay relatively large amounts of money.

In a sense, this is a barrier for smaller ISPs to enter the private peering arena and private peering has evolved as a "good old boys club."[59] The large providers have established private peering with the other large providers and have effectively shut out the smaller providers from the "good" bandwidth. The private interconnects that offer faster transit and less packet loss can certainly be considered "better" bandwidth than the bandwidth at the NAPs, where congestion results in latency and loss.

A smaller ISP could buy transit on discrete circuits from the larger ISPs rather than peering at the NAPs but would then have difficulty competing as a peer due to this added costs. This allows large ISPs to exclude new entrants and smaller players because private interconnects are certainly not offered to all comers and are not really based on any publicly disclosed criteria. Over time, it is quite possible that this private interconnect space will be seen as anti-competitive and ultimately collusive due to these factors.[60]

The End of "Open Peering"

Through 1996, large ISPs generally followed "open peering policies," which meant that they would peer with other ISPs regardless of the other ISP's size, geographical presence, etc. However, in March of 1997, UUNET Technologies announced that it planed to phase out "free of charge" interconnection agreements unless the other networks had substantial national investments in their

[59] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes. May 1998. (p. 6)

[60] Ibid.

own infrastructures.[61] This was further defined as "national networks with dedicated, diversely routed DS-3 networks, 24x7 support staff, and connections to UUNET at four or more diversely located locations.[62]" UUNET notified its peers of the termination date for free peering, stated they would begin to charge as of that date, and forced the other ISPs to sign non-disclosure agreements before the negotiation of such charges could begin.[63] Around this same time, Sprint announced similar plans to not renew "free" peering with certain ISPs.

Both announcements caused a major crisis in the Internet world. There were a great many articles written by newspapers and trade journals chronicling the issue.[64] Many postulated that Worldcom, which had recently bought UUNET, had told UUNET to make the move in order to raise money for Worldcom, which was then suffering from heavy debt.[65] As the ISP business was tremendously competitive, it was difficult to be profitable, and to charge ISPs that did not meet certain criteria for interconnections would mean more potential revenue for all large ISPs. In addition, the dispute between UUNET and Whole Earth Networks, one of the networks that had been notified by UUNET, eventually led to the dismissal of Whole Earth Network's then Vice President David Holub due to Holub's public protests with UUNET[66]. Soon after his dismissal, Holub posted messages to the NANOG mailing list suggesting that such peering cuts were anti-competitive and possibly illegal, and that the government had a role to play to protect public interests by mandating interconnection requirements. This was one of the first instances of

[61] Fleishman, Glenn. "Border Wars on the Net." June 1997. Available at URL http://www.tidbits.com/tb-issues/TidBITS-383.html#lnk4.

[62] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

[63] Cook, Gordon. "UUNET and Sprint Ending Free Peering." The COOK Report on the Internet, Extra Edition. May 3, 1997.

[64] Links to nearly 30 such articles can be found at http://www.clark.net/pub/rbenn/debate.html.

[65] Cook, Gordon. "UUNET and Sprint Ending Free Peering." The COOK Report on the Internet, Extra Edition. May 3, 1997.

[66] Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

proposed government regulation of Internet interconnections. (Holub's arguments will be discussed in more detail in Section III.)

The consequences of UUNET's and Sprint's moves would have made it very difficult for small to mid-size networks to stay in business competitively, as well as made it extremely tough for a start-up to enter the market at all. Such networks would generally have been "non-facilities based,[67]" and therefore needed to lease lines from telephone companies. Having to pay for connections to the largest networks in addition to paying for the leased lines on which their backbones were built would make the prices they had to offer their customers rise to levels far above those of the largest ISPs (which would not have had to pay for peering). Even such large companies like AT&T, GTE, and Intermedia Communications, which all decided to get into the Internet Service business a bit late, had trouble breaking into the free peering "old boys club" of the big five providers (UUNET, MCI, Sprint, PSI and ANS).[68] Instead, they opted to buy CerfNet, BBN Planet, and Digex respectively,[69] all of which had previously established interconnection agreements with the big five.

Recall also that many of the NAPs were owned by ISPs or companies that owned ISPs – Worldcom had purchased MFS, so it owned MAE-East and MAE-West, as well as owning UUNET; Sprint was an ISP and owned the Sprint NAP, etc. These companies were also facilities based corporations. By allowing the NAPs to become congested, the smaller (non- facilities based) ISPs that were there and that had not moved to private peering due to cost issues, were now faced

[67] Facilities-based ISPs have a telecommunications infrastructure of fiber optics, switching equipment, etc, on which they can build their ISP business. An example is a company such as AT&T, a traditional telephone company with an extensive network of fiber optics on which they built an Internet backbone. Non-facilities-based ISPs are those ISPs that have to lease fiber optic lines from telephone companies and build their network on top of those lines. They are at a competitive disadvantage, for several reasons, including the cost of leasing lines, the delays of ordering and receiving such lines, etc.

[68] Cook, Gordon. "Peering and Transit Issues NANOG Discussion Illuminates Some Current Practices and ISP Concerns." The Cook Report on the Internet. July 1996.

with the dilemma of giving their customers spotty service over the NAPs or to pay for interconnections to larger ISPs.

The large ISPs did try to offer some justification for why they had announced the end of free peering for those ISPs that did not meet their newly defined criteria. John Sidgemore, president of UUNET technologies, stated that "UUNET, MCI, and Sprint have invested a lot of money in the [Internet's] infrastructure, and a lot of others haven't. Now we have huge networks and many the little guys want to connect to our network for free. Half our business is reselling our networks to others." [70]

What Sidgemore was describing was the "free ride" that many felt small service providers were after when connecting to larger providers. Take, for example, for a regional service provider that was only connected to one NAP, and wanted to peer with a large ISP. The regional provider would send all of their routes to the large provider, and the large provider would do likewise. The regional provider would get a "free ride" across the large ISPs national infrastructure, as they would exchange traffic at a single NAP, where all of the regional provider's customers were, yet be able to reach all of the large ISPs customers no matter what region they were in. There were proposed technical solutions to this issue, such as the large ISP only announcing regional routes to the regional service provider. However, it was technically difficult at the time and remains so today. It is not impossible, but it is difficult to do and does not scale well with the growth of the Internet.

There was much speculation on the mailing lists NANOG and com-priv[71] about why the large ISPs had announced their intent to stop free interconnection policies. For instance, some thought

---

[69] Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Not yet published.

[70] Dunlap, Charlotte. "UUNET Defends Peer Pricing." CMP TechWeb. May 2, 1997. Available at URL http://www.techweb.com:80/wire/news/may/0502uunet.html.

[71] NANOG is the North American Network Operators group, where operational issues of the Internet are discussed. "com-priv" was a mailing list started to discuss the many issues involved in the commercialization and privatization of the Internet.

that the large ISPs did not want to spend the necessary resources (time, maintenance, network equipment, etc.) needed to establish interconnections with every small ISP that might come along. The benefit for a small ISP when connecting to a large ISP is much greater for the smaller network than for the larger. Another point that was discussed was that the large ISPs were worried the small ISPs didn't have the technical skills necessary to properly configure and manage the interconnection. A misconfigured interconnection between two networks can wreak havoc on one or both networks. However, many others argued that the chance for such issues were just as great when interconnecting with another large ISP.

There was a major backlash from many ISPs, and twelve mid-sized networks, including GeoNet Communications, NetRail, and Whole Earth Networks, banded together to fight UUNET.[72] In order to avoid bad publicity, UUNET grandfathered in many of the complaining networks prior arrangements of free peering, meaning those companies maintained their free of charge peering relationship with UUNET, while new networks that wanted to interconnect with UUNET would have to pay.

Several such incidences have happened since then. In May of 1998, BBN Planet (bought by GTE) told Exodus and AboveNet that they would no longer peer with them. They did not publicly announce this, but it did come into the open on the NANOG mailing list. Again, there were many debates on why BBN Planet would do this, what it would mean to the Internet, etc. Another such incident occurred in October of 1996, when Digex and AGIS cut off their peering connections for over a week, and customers of AGIS were not able to reach web sites at Digex (which included the Securities and Exchange Commission), and vice versa, during that time.[73]

---

[72] Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Not yet published.

[73] Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

Though such incidences have occurred several times since the dissolution of the NSFNet, the Internet has generally remained a fully interconnected network, in which all users had end-to-end connectivity to all other users (withstanding intermittent outages). However, over time, "the standard of openness began to slowly disappear [and] along with it much of the efficiency as well as the relatively even-handed relationships that had allowed market entry based on ingenuity and modest capital."[74]

Anti-competitive Behaviors

In the previous few sections, some possibly anti-competitive behaviors of large service providers were discussed. First, facilities-based companies that owned both NAPs and ISPs had incentives to allow the quality of the NAPs to suffer. By doing this, they were able to segment Internet interconnections into good bandwidth (private leased lines) and bad bandwidth (the public NAPs), in what could be viewed as an attempt to coerce small, non-facilities based ISPs to use private peering instead of public peering. Since in most such private interconnection arrangements, each ISP pays for half of the private circuits, and because the facilities based ISP owns the circuits, it is clear that the facilities based ISP gains a financial advantage. Second, several large ISPs tried to move away from open peering, to a model in which they charged smaller ISPs for such interconnections. Again, this move gave a financial advantage to the large ISPs, forcing smaller ISPs to raise prices to the point where they could not compete with the large ISPs. It also made market entry by new comers extremely difficult because they could not freely interconnect with other ISPs, but instead had to pay for such connections.

[74] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

The large service providers know that if they don't have connectivity to small ISPs, that the large ISPs' customers will not notice and/or care as much as the small ISPs' customers would. The small ISP has less to offer the large ISP in terms of the number of hosts the large ISPs' customers would want to access. On the other hand, the small ISPs service would suffer severely without connectivity to the large ISP, as its customers would not be able to reach large sections of the Internet. In order for the small ISP to keep their customer base intact, they must be able to reach the large ISPs' customers. In this sense, the large ISP has the upper hand in forcing the small ISP to pay for interconnections to the large ISP. The small ISP could also chose to pay for connections to another ISP that does have connections to the large ISP, but, in either case, they are still paying for the connections that at one time were offered free of charge.

In the case of UUNET's announcement in May of 1997, they had announced the criteria they would use in determining whether or not they would peer with other providers. The criteria included a national infrastructure of DS3 size and the ability to interconnect with UUNET at four geographically distinct locations. Other large ISPs had announced similar requirements since then. Yet many mid-sized ISPs had willingly built out their networks to meet such criteria, only to find that the large networks still would not peer with them for one reason or another (which was never clearly explained).[75] There were also complaints that in some cases the peering criteria were not made publicly available to begin with.[76]

Government regulation occurs when there is potential for and/or evidence of anti-competitive behaviors. In 1890, the Congress of the United States passed the Sherman Antitrust Act, which declared illegal "every contract, combination in the form of trust or otherwise, or conspiracy, in restraint of trade or commerce among the several States, or with foreign nations." In 1914, the

[75] Cook, Gordon. "David Holub Explains Model of an Open Peering, Open Interconnect, 'Bill and Keep' Internet." The Cook Report on the Internet. July 1997.

[76] Ibid.

Sherman Act was amended by the Clayton Antitrust Act "to prohibit certain monopolistic practices that were then common in finance, industry, and trade.[77]"

The original interconnection model of the Internet, which was "characterized by widespread, voluntary and non-discriminatory interconnection," has begun to shift to a "hierarchical discriminatory model[78]" in which large ISPs do not freely interconnect with all other ISPs. This shift could lead to a "balkanization" of the Internet, in so far as there is a "big boys club[79]" of large ISPs that are possibly inaccessible from non-member networks unless there is compensation from that network to one that is a member of the club. The "big boys club" referred to is a group of several large ISPs that together have a large majority of the Internet users, and which is seen by some as possibly acting in collusion to force smaller ISPs to pay for interconnection to the club.

A trend of recent years has been the consolidation of ISPs, thus forming larger and larger networks. The Internet has remained competitive due to the number of companies providing service, but the system is fragile and if a single market leader emerged that was much larger than all the others, the stability could be in trouble.[80] Today, if a single network refused to interconnect with others for any significant length of time, it would likely put itself out of business because all of its customers would lose end-to-end connectivity with the rest of the Internet, and would therefore migrate to other providers. However, if one company were able to gain enough market share such that other companies could not provide a useful service without connecting to the "giant" network, then that company could potentially use that dominate position to either take over the market or extract payments from the smaller companies. "A network system tends towards monopolization

---

[77] MSN Encarta Online. Available at URL http://encarta.msn.com.

[78] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

[79] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes. May 1998. (p. 6)

whenever one company has enough of the market to conduct an effective business without interconnection with other companies.[81]"

An example of a merger that could have had such an effect was the merger of MCI and Worldcom (which owns UUNET). At the time, Internet MCI and UUNET were the two largest ISPs, and had they been allowed to merge, could possibly have dictated peering relationships. Both the FCC and the EU ruled that the combined entity would control too much of the Internet market and could therefore be potentially damaging to a competitive environment. Worldcom had also purchased CompuServe and ANS (along with several other smaller providers) and those, with the addition of Internet MCI, were seen as creating too dominant of a player in the market. In the end, the Worldcom merger with MCI was approved with the stipulation that Internet MCI be divested. Internet MCI was sold to Cable and Wireless for 1.75 billion dollars and the rest of the MCI – Worldcom merger moved on.

In addition to a trend of consolidation of ISPs creating potential competitive imbalances, there is also a "concentration of control of Internet facilities" in which national ISPs control a majority of the NAPs. The NAPs were first awarded to Sprint, Ameritech, Pacific Bell, and MFS, and at that time, only Sprint had a significant ISP business. Therefore, the neutrality of the NAPs was fairly strong. However, the neutrality of the NAPs is now in doubt, as Worldcom bought MFS (which ran the MAE's) and UUNET, so the worlds largest ISP and the most heavily used NAPs are now owned by one company. In addition, Pacific Bell has a strong Internet presence, as does Ameritech. Southwestern Bell has already purchased Pacific Bell, and is in the process of merging with Ameritech; therefore one company that provides significant Internet services will also own both the Pacific Bell and AADS NAPs. When large ISPs own public NAPs, there is they have

---

[80] Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

[81] Brock, Gerald W. 1994. "Telecommunications Policy for the Information Age." Cambridge, Mass. Harvard University Press.

incentive to allow the quality of the NAPs to suffer so that smaller ISPs are forced to privately peer and/or buy transit service. Such concentration of control has also led to a hierarchy among the ISPs, with the large ISPs at the top of the pyramid using their market power to drive the small ISPs out of business.[82]

These developments and actions – namely the hierarchical discriminatory model with a concentration of control of Internet facilities in which the large providers have the potential to force smaller providers into paying for interconnections -- could be said to fall under the Sherman and Clayton Antitrust Act, as defined above.

The move away from such open interconnection, in which no money changes hands ("settlement free"), to ones where certain ISPs are trying to charge for such connections, closely resembles the interconnection agreements that are in place in the traditional telecommunications carrier-to-carrier connections.[83] The next section will closely examine those agreements and compare them to what is currently in place on the Internet, and highlight why it is difficult to use such models for the Internet.

---

[82] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

[83] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

Section II: Telephony and Internet Interprovider Settlements

Section I explored the history of the Internet and examined how Internet interconnections between commercial entities came into being when the single Internet backbone, the NSFNet, was dissolved, and a multiple backbone architecture emerged. It was noted that no real consensus was established on inter-provider interconnections, i.e., which networks should connect to each other, which should not, and what the terms of those connections should be. This section will examine interconnection and settlement agreements for both the traditional telephony world and the current (mostly) "settlement-free" agreements in use for Internet interconnections. Telephony arrangements are first discussed as they have evolved over a much longer time and provide background for the Internet arrangements.

Traditional Telephony Interprovider Settlements

In the past, telecommunications providers generally interconnected to one another voluntarily as they sought to increase their geographical coverage as well as their network reach.[84] However, with the introduction of competition in both the long distance and local telephone service provider market, legislative and regulatory rules that mandate carrier-to-carrier interconnection agreements became necessary. That is because the incumbents generally do not need to interconnect to new comers to promote geographical expansion or network reach, and in fact have revenue incentives not to interconnect to competition in most cases. When there is competition, the incumbent can lose

[84] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998. (p. 7)

market share and therefore lose revenue. To help foster competition, regulations are often needed to force the incumbents to interconnect with any upstarts.

Traditional Telephony Interconnection Requirements

In the United States, telecommunications companies, as mandated by the Communications Act of 1934 (section 251a) must "interconnect directly or indirectly with the facilities of other telecommunications carriers."[85] This requirement is basically one of universal accessibility, meaning that all communications providers are interconnected with one another in some manner such that one provider's customers can get to any other provider's customers. (This is the same premise of end-to-end connectivity found on the Internet.) It should be noted that at this time there is no such requirement for Internet Service Providers to interconnect, and, in theory, the lack of such a requirement means that the Internet could segment in such a way that not all users had universal accessibility to all others. However, in practice, this has not yet happened, mostly due to market pressures and the demands of the installed customer base of the ISPs.

Traditional Telephony Settlements

In addition to being required to interconnect by federal mandates, common carriers have financial "settlements" between them. The most general kind of such arrangements is often termed "bilateral settlements.[86]" In this scheme, a "call-minute" is the unit used for accounting.

---

[85] Communications Act of 1934, as amend but the Telecommunications Act of 1996. Available at URL http://www.fcc.gov/Reports/1934new.pdf. (p. 91)

[86] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p. 5)

Internationally, when a client originates a call, the client's local service provider charges them for the duration of the entire call. The call may pass through (transit) one or more providers before being terminated on the remote client's local network. Most such transactions involve just two providers, the originating and terminating ones. In this simple case, the terminating provider charges the originating provider a per-minute cost with an accounting rate that has been pre-negotiated between the two parties. Both parties can charge each other depending on which one terminates a given call. At the end of the accounting period, there are two sets of call-minutes, one in each direction, and when there is an imbalance, the provider hosting the great number of originating call minutes pays the other according to the negotiated rate. In more complex cases where more than two providers are involved in the call, the settlement happens on a hop-by-hop basis between the two networks at each hop.

Another type of settlement scheme used in international telephony interconnections, though rare, is "sender keeps all (SKA)."[87] In this case, each service provider invoices its own clients for the end-to-end service they are provide, but no financial settlement is made across the interconnection structure. In a sense, SKA can be considered a boundary case of the bilateral settlements previously described, in which both parties decide the outcome of the call accounting process to be equal whether or not that is actually the case. Because no accounting takes place and it is assumed balanced, no financial settlement is payable by either party.

Another type of settlement used in the International market is "transit," in which one party invoices the other for services provided.[88] In this case, one network pays the other to provide connectivity to the rest of the users on the network, whether they are on the other network or some network farther down stream. This is commonly used as the basis for long distance provider to local

[87] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p. 6)

[88] Ibid.

access provider interconnections. In a sense, this is also a boundary condition of the bilateral settlement model, in which the parties agree to pay in only one direction.

Domestic Regulations

In the United States, there are two main types of telephony settlements in place – access charges and reciprocal compensation, both of which are regulated by the FCC via sections 251 and 252 of the Communications Act of 1934, and further amended by the Telecommunications Act of 1996.[89] Both of these settlements are specific cases of the bilateral settlements as described above. In addition to these settlement requirements, the Acts also mandate contributions to the Universal Service Fund (section 254d). All three of these regulations warrant discussion here, as they are currently being challenged due to the emergence of Internet access providers. These will be discussed in more detail in section III, The Internet and Current Regulation Issues.

Access Charges

Access-charges have to do with long distance calls. For each long distance call a consumer makes, the long distance carrier must pay per-minute fees to the originating and terminating local telephone service providers. These fees are designed to recover the costs of the local telephone company for the use of their facilities, and include both outgoing and incoming calls.[90]

---

[89] Telecommunications Act of 1996. URL: http://www.fcc.gov/telecom.htm.

[90] FCC Access Charge Reform Homepage. 7/16/99. URL http://www.fcc.gov/isp.html.

Reciprocal Compensation

Reciprocal has to do with local telephone calls. In this case, a local telephone company is compensated for completing a local call that is placed by one of its competitor's customers[91]. Typically, the companies involved are an incumbent local exchange carrier (ILEC), which often serves a large number of customers, and a competing local exchange carrier (CLEC) that has just recently entered the market and therefore has fewer subscribers.[92]

Universal Service Fund

The Universal Service Fund was established by the Communications Act of 1934 to subsidize basic telecommunications services for low-income individuals and the costs associated with bringing service to high cost rural areas[93]. The cost of providing service in urban areas is essentially much cheaper than in rural areas due to economies of scale. (I.e. when creating the infrastructure to support telecommunications, less money has to be spent when there are a large number of customers in a small area than when there is a large area and a small number of customers.) The telephone company can service urbanites at cheaper rates than what they actually charge. The extra monies are used to subsidize the rural users, who pay less than what it actually costs for the telephone company to provide them service. The FCC deemed this necessary in 1934, and still does, so that basic telephone service can be available at low prices to all households in the United States.

---

[91] FCC Reciprocal Compensation Fact Sheet. February 1999. URL: http://www.fcc.gov/Bureaus/Common_Carrier/Factsheets/nominute.html.

[92] "Reciprocal Compensation Fact Sheet" URL: http://www.fcc.gov/Bureaus/Common_Carrier/Factsheets/nominute.html.

[93] Geist J., Rudolf. "Policy Forum: Regulatoin of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

Section 254(d) of the 1934 Communications Act, as amended by the Telecommunications Act of 1996, mandates universal service contributions from "every telecommunications carrier that provides interstate telecommunications services." The 1996 Act went further and expanded universal service such that specific beneficiaries -- namely schools, libraries, and rural health care providers -- receive federal funding to subsidize Internet access.

Telephony Settlement Trends

Both the international and domestic telephony models are by no means stable. In the United States, for instance, there are court cases and FCC hearings on both access charges and reciprocal compensation due to the emergence of the Internet and the access issues it has brought forth. (These issues will be discussed in more detail in Section III.) On the international front, there is lots of pressure to move away from bilaterally negotiated uniform rates to separately negotiated rates in each direction.[94] In addition, deregulation is changing the transit fee model as local providers extend their networks into the long distance market and begin interconnection agreements with companies in similar positions. In both domestic and international cases, there is much criticism because the negotiated rates often are not cost-based, but are instead based on a desire to create revenue stream in the direction of the provider with more "power."

Current Internet Interconnection Agreements

Recall from Section I that at the time the NSF backbone was dissolved, there were no guidelines given for when ISPs should interconnect or under what terms such interconnections

[94] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p. 6)

should be made. In those days, however, network operators favored network connectivity over desires to make short-term profits. In other words, most networks were open to connect to all other networks just to maintain end-to-end reachability of the Internet, whether or not the interconnections made financial or architectural sense. ISPs "readily accepted and routed traffic generated by other [ISPs} with little regard for the balance of traffic flows, network size/ geographical reach, etc."[95]

That is not to say that it wasn't debated at all – it certainly was, and quite vocally in the Internet mailing list "com-priv," which was originally established to discuss the commercialization and privatization of the Internet. The debate still shows up several times a year on both the "com-priv" and "NANOG" mailing lists. Both of these mailing lists have in the past and today continue to shape the Internet, yet no conclusions or decisions have been drawn in all the debates. The discussions have centered around which networks should interconnect as peers (i.e. there is no customer – provider relationship and no money passes between the two), when there should be a customer-provider relationship (transit), etc.

Transit Agreements

Transit agreements are established when there is a clear customer/provider relationship between two networks. In such cases, a small or regional network needs to get access to the rest of the Internet, so they pay a large provider for "transit" services. In such cases, the large provider sends either a default route[96] or the entire Internet routing table to the customer network. In addition, the large network advertises the customer's routes to the rest of the Internet. These route

[95] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

[96] A default route is also called a gateway of last resort, and means that any packet received that has a destination not in the route table, should be sent to the gateway of last resort.

announcements and traffic exchange enable the small network's customers to reach both the large network's customers as well as the rest of the Internet. The smaller network pays the larger network for this service based on bilaterally negotiated terms.

Peering Agreements: Sender Keep All (SKA)

Peering agreements are established between two networks when "both parties perceive equal benefit from the interconnection." [97] In such a relationship, each ISP sends their own routes and those of their customers to the other ISP. (Note how this differs from a transit relationship, in which a default route or the entire Internet routing table is sent. By sending only their own and their customer routes, there is no connectivity provided to "rest of the Internet," and that connectivity must be obtained via additional interconnection relationships – either peering or transit.)

The financial arrangement used in such peering relationship is generally "bill and keep/sender keep all." This means that the exchange of traffic between two ISPs occurs with no mutual charge – i.e. it is an interconnection agreement without financial settlements. Each ISP retains all of their customer payments in full without having to settle accounts with other ISPs who participate in the routing and delivering of that customer's traffic. This method works well when there is nearly symmetrical traffic flow[98] and the cost of terminating the traffic is low compared to the cost of

[97] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p. 13)

[98] Ibid.

metering it.[99] This model has served as the primary arrangement between "peer" providers due to "the administrative convenience and willingness of ISPs to promote network connectivity."[100]

SKA works well and is most stable when both parties involved perceive equal benefit from the interconnection.[101] In this sense, the providers must be of "equal dimension" in geographic scope, customer base, traffic flow, market power, etc. – namely whatever each party perceives as important. Each provider makes an independent assessment of the value of the interconnection, and if each provider independently concludes that an interconnection will provide some mutual benefit, and interconnection is made will generally remain on a stable basis. However, if one party begins to perceive that the relationship is no longer balanced, and that the SKA interconnection is resulting in the leverage of their infrastructure by the other ISP, the relationship becomes unstable. (Recall the "free-ride" described by John Sidgemore, President of UUNET technologies in "The End of Open Peering section." Also recall from that section when UUNET, Sprint, and others announced the end of "free peering" (a.k.a. SKA). It was because they began to perceive that there was no longer balance in some of their peering relationships.)

It should be stressed that the essential criteria for a stable SKA relationship is the perceived equality that the networks achieve by interconnecting on a settlement free basis. This is a subjective evaluation that each ISP must make for each established and potential interconnection relationship. Because it is subjective, each ISP is likely to come up with their own set of peering criteria that they use to evaluate other ISPs before entering a peering agreement. Since no money changes hands, if either IPS feels there is an imbalance in anyway, they may feel that the other network is cheating them. Due to this, some have suggested that money should flow towards the ISP that is providing

---

[99] Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

[100] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

[101] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p. 13)

greater value – i.e., they want to establish settlement-based interconnections. However, because the determination of relative value is strictly subjective, determining which direction the money should flow is a controversial topic. In addition, there are many issues with determining a fair settlement mechanism on the Internet, as described in the next section.

Problems with Settlements on the Internet: What could be the "call-minute?"

Over time, ISPs have appeared more inclined to interconnect facilities only if a transfer of payment occurs, and slowly the Internet is becoming more like a system of telephone company networks.[102] However, there are many problems with treating Internet interconnections and telephony interconnections similarly. Perhaps the biggest stumbling block is determining a fair accounting unit. With telephony, a call minute is the standard unit and is simple to account for; however, on the Internet there is no such simple accounting unit to available to use. This section explores proposed several ideas and their associated problems.

Packets?

One proposed equivalent to the "call-minutes" used in telephony settlement arrangements is TCP/IP packets. Packets are at the heart of all Internet traffic – everything that crosses the Internet does so as packets. Recall that with telephony settlements, the terminating provider charges the originating provider a per-minute cost with a prescribed accounting rate. Therefore, with packets as the unit of accounting, the network receiving the packet would charge the network that originated the packet. Packets flow both ways in each individual session that crosses the interconnection

[102] Ibid.

boundary. At the end of the accounting period, the network that received more packets than it sent would receive a settlement payment from the other network.

There are many problems associated with using packets as the call-minute. First, it is extremely difficult to keep precise measurements of packets in the core of the Internet due to the tremendous volumes of traffic. Current Internet technology, which includes hardware such as routers and switches and the software that runs those, has not been designed for such accounting purpose and therefore has difficulty doing it without causing resource strains on the equipment. Also network outages, data collection failures, etc., can make it difficult to keep accurate records. Second, this method is open to abuse as bogus traffic can be created in either direction to skew the balance of traffic flow, and such bogus traffic is very difficult for either party to notice.

Session flows?

Another proposed accounting unit is for session flows. Each TCP session can be considered to have two flows, one in each direction. Certain router vendors already have "flow" data gatherers, such as Cisco's NetFlow[103]. By measuring session flows in this manner, the actual settlement structure could be based on the duration of the flow (similar to telephony's call minutes) or traffic volume of the flow.

There are also issues with this method. First, not all traffic on the Internet is part of a TCP flow. There is also the UDP protocol, which has no real session characteristics., as there is no acknowledgement from the receiver to the originator, etc. However, for UDP traffic, there could be some type of accounting against the source of the packets. There is also the same issue of

[103] NetFlow Services and Applications. Cisco Connection Online (CCO). August, 1999. http://www.cisco.com/warp/public/cc/cisco/mkt/ios/netflow/tech/napps_wp.htm.

technology that was discussed with packet accounting – namely that the current Internet technology was not built for such accounting, and that it is extremely difficult to do such accounting in the core of the Internet due to the tremendous number of concurrent sessions over a single exchange point at any given time. Finally, there are issues with asymmetric routing, in which it would not be possible to see both sides of the flow from one exchange point.

Traffic volumes?

Another proposed accounting unit is traffic volume. In this scenario, the volume of traffic passed in each direction across the interconnection point would be measured, and at the end of each accounting period, the two ISPs would financially settle with each other at their agreed upon rate as applied to the net traffic flow. However, which way the money would flow is not entirely obvious. One possible model is that the origination network should be funding the terminating provider to deliver the traffic, and therefore the money flows in the same direction as the traffic. However, another model assumes that the overall majority of traffic is generated in response to actions of the receiver (i.e. a web page request, a request to transfer a file, etc.). In this case, the terminating provider would pay the originating provider. (These two different views will be explored in more detail in the next section, "Defining Who has Value.")

Accounting based on traffic volumes has some of the same problems as packet counting and flow accounting. In the case of traffic volume, it is much easier to measure than the others (and in fact is currently done by most every provider on their interconnections to measure utilization for capacity planning reasons). However, this is also open to abuse, as it is easy to create bogus traffic to shift the revenue to a given network. "Overt abuse is often easy to detect, but greed is a wonderful stimulate to ingenuity and more subtle forms of abuse could follow."[104] Also, third parties could create abuse with such tactics as loose source routing[105] and source address spoofing.[106]

Defining who has Value

As seen above, the view different people take on which network provides value to the other is not always clear. For any two ISPs that interconnect with each other, the decision as to which one should assume the role of upstream provider and which should assume the role of downstream provider is not always easy to determine for either party, or even to an outsider for that matter. The question becomes "can an objective determination be made of whether an ISP is a peer to or a client of another ISP?" In many cases, it appears that the question is not even appropriate because the role of client and provider can change in various instances. However, it is a critical question, because if a completely objective determination cannot be made, then question becomes "who is responsible for making a subjective determination, and on what basis?"[107]

It is interesting to note that such value issues were discussed as early as March of 1990, at the conference held for discussions of the "Commercialization of the Internet," as documented in RFC 1192, "Commercialization of the Internet Summary Report."[108] At that conference, it was noted

---

[104] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj

[105] Source routing allows a sender of an IP packet to dictate the path the packet takes through the network, rather than letting the normally used routing tables dictate the path.

[106] Spoofing an address can loosely be described as forging the address. This method is often used in Denial of Service (DOS) attacks that generate bogus traffic to the victim.

[107] Huston, Geoff. "Interconnection, Peering and Settlements – Part I." Internet Protocol Journal, Volume2, Number 1, March 1999. Soft copy available at URL http://www.cisco.com/ipj

[108] Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

that the large networks of the time, namely CSNET, BITNET, UUNET, PSI etc., were using the NSFNet backbone for free, so it appeared they were enjoying a subsidy from the NSF, Merit, IBM, and MCI, the four entities that were funding the NSFNet. However, they generally added their own value to the nascent Internet as all users benefited from access to each other's users and resources. It was also noted that many small or startup networks generally brought in less resources, so they may have benefited more from their connections than the other parties. And finally, the conference touched on traffic asymmetry issues. For example, three times more traffic flowed from the Internet to MCIMAIL than from MCIMAIL to the Internet. At the time, there was a sender-pays fee structure to MCIMAIL such that the sender had to pay for mail to be delivered. Therefore, the network originating more mail had to pay the other network. However, today, there are no settlements for peering, and therefore money does not transfer between the two peering networks.. Some feel that such large imbalances in traffic flow are inherently not fair if not money changes hands.

The Internet has evolved with three main types of ISPs: those that are dial-centric, meaning most of their customers use the ISP as a dial connection to the Internet, those that are Web-centric, meaning most of their customers host content on web pages on the Internet, and those that are business oriented, in which the majority of their customers are businesses that purchase leased lines from the ISP to access the Internet. The client-server architecture of the Internet leads to asymmetric traffic flows. When a user requests a web page, they send a relatively small packet to the web page's server, and that server sends the entire web page back to the requester generally in many larger packets. As multi-media rich web pages have become the standard web fair, large amounts of data are sent back to the requester. The same is true for file transfers. A user sends a few small packets to request what may be very large file. The resulting imbalance in traffic flows for such sessions can be anywhere from 3 to 1 to 15 to 1. This asymmetry is beyond anything reasonable, and with every new generation of Internet applications, from email, to FTP, to WWW, to streaming media (such as RealAudio/RealVideo), the traffic asymmetries from requester to sender have become larger.

The asymmetries can sometimes be in the reverse direction. For example, an email sender is the session requester, and if they have attached large files to the email, then the requester is sending more data to the receiver than the receiver to the sender. This is the opposite of what happens for the common Internet applications of HTTP (web browsing), FTP (file transfer), and streaming media, but does show that no assumptions about traffic symmetry are valid for all types of traffic that can be found on the Internet.

So the question of who provides value to whom is controversial. Is the value provided from the web-centric ISP for having all the content, or is the value provided by the dial-centric ISP for having the users that want to access the content. It is also apparent that measurements such as greater customer base or greater geographic coverage have little to no meaning. For instance, if one ISP has 10,000 users, and one has ten, which one would be expected to gain the most by interconnection to the other? If those 10,000 users of the first ISP are all dial users, and those ten customers of the second ISP are the largest ten Web sites on the Internet, is it even possible to define?

The costs associated with dealing with a content centric service provider can be quite significant. The Internet evolved with “hot potato” routing, or “shortest exit routing.” This means that traffic from network A to network B is sent to network B at A’s closest connection to B, no matter where the final destination in network B may be. The Internet evolved this way, in part, because it saves the sending ISP costs in terms of network infrastructure. What this means for content-centric providers is that wherever their data centers are, they dump traffic to the requester at the point closest to the data center that they can. As the session flow is asymmetrical, in the sense that the requestor sends very little traffic to the server while the server sends lots of traffic back to the requester (for most types of traffic), it is the requester's network that needs the large infrastructure to handle the session. In such cases, "open" or "settlement free" peering often benefits the web-centric provider the most in terms of the cost of network infrastructure build out.

There are ways to do "cold potato" or "longest exit" routing, in which the content centric provider would carry the traffic on their own network to the interconnection point nearest the requestor, but there are some issues with this. Longest exit routing is accomplished by setting the BGP[109] knob Multi Exit Discriminator (MED) on route announcements. For example, if network A peers with network B on both the East and West coasts, network A can mark all routes that are "near" the East coast with a low MED on the interconnection point on the East coast, and mark those same East coast routes with a high MED on the West coast interconnection point. In this way, network A is able to give network B some information about where the final destination is on network A. See Diagram 8 for an example.

[109] BGP is the routing protocol used between service providers.

Diagram 8: Network A peers with Network B on both the East and West coasts. The route prefix x.y.z.0/24 is connected to Network A's San Francisco node, so network A advertises that route to Network B with a lower MED on the West coast than on the East Coast. Network B will choose the route with the lowest MED to send packets to network A.

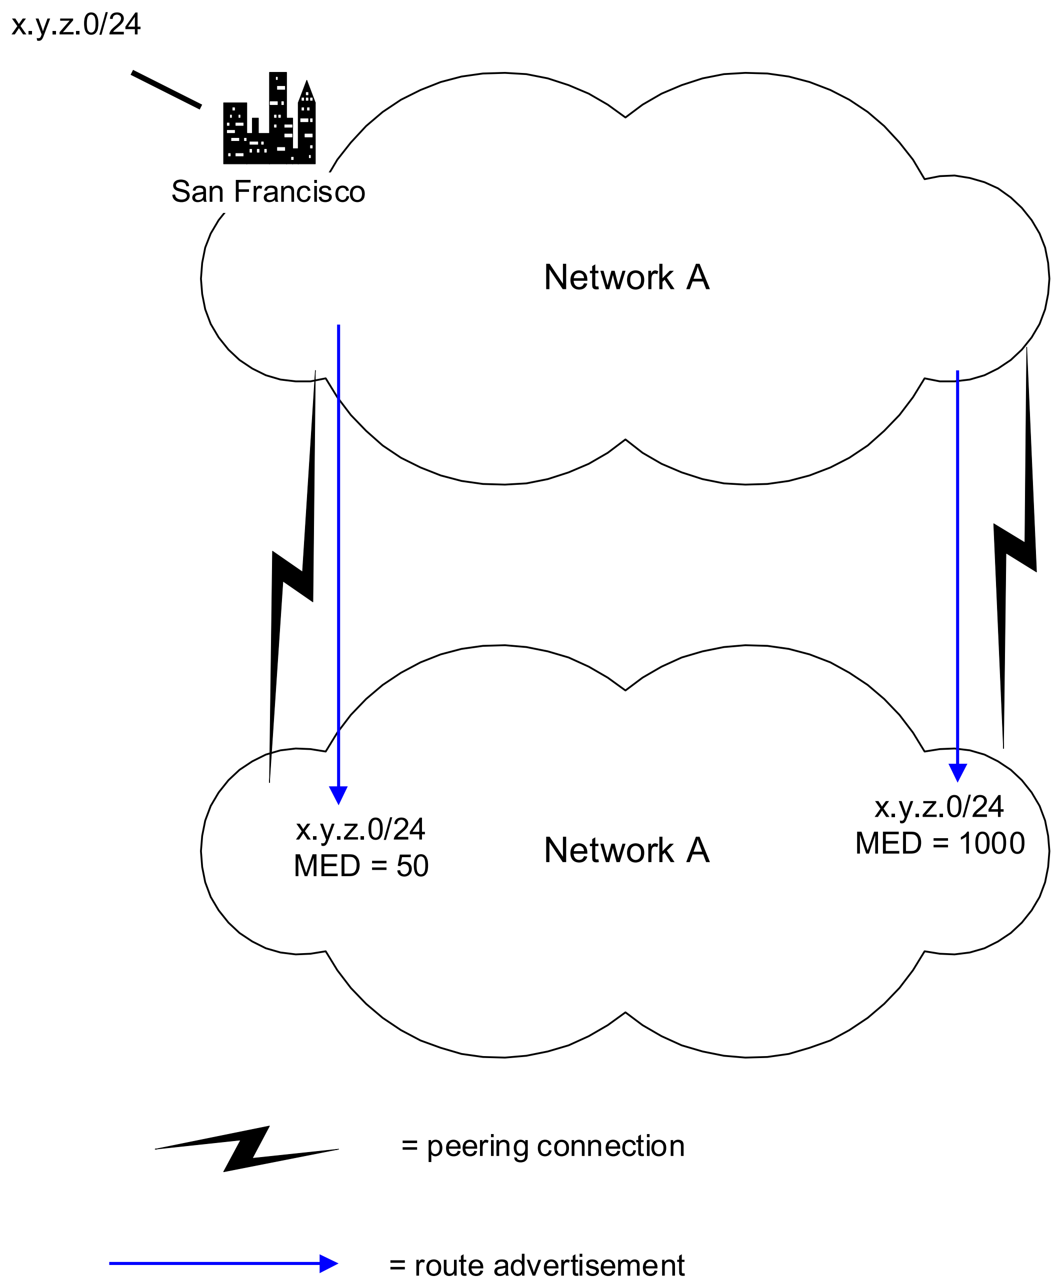

x.y.z.0/24
San Francisco
Network A
x.y.z.0/24
MED = 50
Network A
x.y.z.0/24
MED = 1000
= peering connection
= route advertisement

One problem with this method is that as the Internet routing tables were growing at an exponential rate, route aggregation became necessary to keep the routers stable. Aggregation means taking networks that are numerically next to each other and announcing them as one route. For example, if a given provider owns traditional class C networks 200.100.1.0 255.255.255.0 and 200.100.2.0 255.255.255.0, which are expressed in Classless Inter-Domain Routing [CIDR] format as 200.100.1.0/24 and 200.100.2.0/24, they can announce one route that covers those two: 200.100.1.0/23.[110] Often, as ISPs evolved, address allocation was not done in a hierarchical or geographical fashion. In the example above, the ISP may have used 200.100.1.0/24 on the East Coast and 200.100.2.0/24 on the West Coast. Therefore, if they are now advertising 200.100.1.0/23, assigning a MED to differentiate approximate location makes no sense. The only solution would be to send the more specific /24's, which eliminates the desired benefits of CIDR – namely reducing the size of the Internet route table. Therefore, that solution does not scale well, as the Internet route tables could grow too large for the routers to handle.

It should be noted that some content-centric providers, such as AboveNet, Frontier Global Center, etc. will listen to MEDs to help bear the cost of the network infrastructure needed to serve their customer's content to the requestors. But not all networks are able to provide meaningful MEDs per above, and again, this may not scale over the long run.

Internet Retail Cost Structure as an Impediment to Settlements

In addition to the problems of determining an acceptable accounting unit that could be used for Internet interconnection settlements, there is another inherent problem because of the many different Internet retail cost structures that different ISPs incorporate. The various pricing plans

[110] Please se the CIDR FAQ for more information, which is available at URL http://www.rain.net/faqs/cidr.faq.html

that ISPs incorporate in billing their customers makes it extremely difficult for the ISPs to uniformly settle with each other at interconnection points.

The Internet began with a "flat rate, all you can eat" cost structure.[111] What this means is that for a given monthly fee, a customer is able to use the Internet as much as they want. Such a pricing scheme makes sense during promotional periods in a nascent market when service providers are trying to garner market share and have adequate capacity to handle the loads. During such times, prospective customers require such incentives to stimulate their interests enough to make the up-front investment in a personal computer and an Internet access subscription. In the early days of the network, there were fewer customers connecting to the Internet, and partially because of that there was some excess capacity. In such times, the cost of adding another customer (i.e. carrying a bit more traffic) is basically negligible.[112]

However, over time, the billing structures have changed to reflect a growing market, and now there is a large variety of billing methodologies. While dial up customers long enjoyed subscription fee's of $19.95 for unlimited access, several of the large ISPs such as AOL, the IBM Global Network and AT&T WorldNet have now begun to charge $19.95 for the first one hundred hours of access, but have additional per hour usage fees beyond that. Besides the dial-up users, business connections, ISP resale connections, and content hosting connections have evolved from flat rate pricing to various different billing models. Some are now charged on sent volume, some on received volume, some on a mix of sent and received volume, and some on the access capacity irrespective on link utilization.[113]

---

[111] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

[112] Ibid.

[113] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj

Given that the end user is paying their ISP for comprehensive Internet connectivity, when a packet passes over an interconnection point, what tariff model should be followed? If the sender's ISP uses a flat rate model for their customers, but the receiver's ISP uses a usage-based model, which model should the interconnection point use? "The financial structure of interconnection must be an abstraction of the retail models used by the two ISPs."[114] Yet because of the heterogeneous nature of the retail models, such an abstraction is nearly impossible. If there was a uniform retail model in which the sender of traffic pays, then client sending traffic pays the first ISP, and that ISP pays the second ISP a tariff to deliver the packet to the destination. On the other hand, if the uniform retail model is one in which the receiver pays, then the second ISP pays the first ISP. But because each ISP may be billing their customer in different way, there is no uniform retail model to apply to the interconnection.

Settlement Costs

Another important aspect of settlements is the cost associated with them. If an acceptable accounting unit were found, it would have to be measured closely. As already stated, the current technology in use has not been designed to do such accounting. However, if we assume that a future generation of technology is developed that can, the costs of doing so can be analyzed. The first cost is in the technology itself, as additional features require additional funds. However, the costs associated with the core technology, which will be doing the normal routing of packets, etc., may be minimal when compared to the cost of the equipment needed store, process, and analyze the accounting data.

At some point, the overhead of such accounting may not justify the process of settlement itself. For example, if it costs network A $10 to determine that it owes network B $100, and it costs

[114] Ibid.

network B $10 to determine that it owes network A $105, network B would need to pay network A $5. Yet it took network A and network B $20 to figure that out. In the telephony market, these charges were tiny compared to what the customer paid in per minute charges.[115] However, a recent RAND study concluded that 90% of the cost of telecommunications goes towards billing and accounting (i.e. not just the billing department, but the network architecture that permits the billing)[116]. The telephony network evolved over many years to incorporate the billing mechanisms, but the Internet has evolved over a relatively short period of time with no such mechanisms included, and at this point in its development, it could potentially cost more money in terms of technology to do such billing than the money that would be generated if it were actually done.

Before moving on to look at the current interconnection arrangements of the Internet, one further point bears discussion. In the telephony world, the "calls" are very homogenous to the network, in that a circuit of 64kbps is established between the caller and the receiver for the duration of the call. Because of this, it is difficult, if not impossible, for one provider to try to fraudulently adjust the accounting scales in their favor. However, on the Internet the traffic is very heterogeneous, in that whatever accounting unit was decided upon for settlements, it would be very difficult for providers to monitor the traffic to verify that it was all legitimate. Bogus traffic that would shift the revenues in one direction or another could easily be generated yet would be difficult to detect and trace. In addition, when an ISP is richly interconnected with many others ISPs and there are settlements involved in those interconnections, the primary objective of a routing policy no

[115] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes, May 1998.

[116] (R24) Cannon, Robert. "Legal Developments: FCC releases Reciprocal Compensation Order. Boardwatch Magazine. May 1999.

longer would be to support optimal network connectivity, but instead would be policy with a primary objective of maximizing financial returns to the network.[117]

Comparing and Contrasting Telephony versus Internet Settlements

At this point the major similarities and differences between traditional telephony interconnections and settlements and those of Internet Service Providers will be reviewed. One major difference is that telecommunications carriers are required by government regulation to interconnect with one another as a way to ensure universal access to all customers of all providers. For ISPs, there is no such requirement mandated by government regulations, and thus far market pressures have forced ISPs to remain interconnected to ensure such universal end-to-end connectivity. There could be a time when the market is not able to force this, at which time the government may need to step in for the good of the public.

Another difference is that telecommunications carriers have financial settlements between them, based on the accounting unit of "call minutes." With traditional voice services, a voice circuit is established from end-to-end and the duration of the call is measured. The network that terminates the call charges the network that originates the call, and because calls generally flow in both directions, the net difference in call minutes is determined with the network that terminates more call minutes receiving payment from the network that originates the calls. For ISPs, transit relationships involve one ISP paying another for carrying their customers traffic to the Internet and is generally used when there is a clear customer-provider relationship. Peering relationships, on the other hand, are between ISPs of "similar size" (as perceived by each ISP) and there are no financial settlements. Settlements are troublesome for the Internet for several reasons, including the non-

[117] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj

existence of a usable accounting unit (e.g. the call-minute of carrier-to-carrier interconnections), the fact that there has been no objective way to determine who provides more value in the relationship, and the fact that Internet traffic flows are much more heterogeneous than telephony calls which makes it easier to fraudulently shift the accounting in one networks favor.

Telephony financial settlements derived from a different set of initial premises than those that fit the Internet, as described by Geoff Huston and paraphrased here.[118] The typical starting point for carrier-to-carrier settlements is that the retail offering of the provider is one of a comprehensive, end-to-end service, in which the originating service provider utilizes the services of other providers to complete the delivery of the service. The originating provider then settles with those other providers who have taken some kind of role in providing the service. With such cost-distribution, both small and large providers are able to operate with some degree of financial stability, which in turn allows a competitive market to exist. However, the retail model of the Internet is not necessarily one of end-to-end service, but one of partial path service. And there is no price component that covers the complete path. "In an abstract sense, the Internet can be viewed as a collection of bilaterally funded path pairs, where the sender funds the initial path component and the receiver funds the second terminating component."[119]

It is interesting to note that Internet interconnections of the non-transit relationship type have begun to raise the same questions of access, equity, and pricing of local and inter-exchange carrier interconnections. However, ISPs have operated under what the FCC terms "enhanced service provider status" and have therefore not been subject to regulation. Such regulatory issues are covered in the next section.

---

[118] Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj (p.19)

[119] Ibid.

Section III: Internet Interconnections and Regulatory Issues

Section I of this thesis explored the history and development of Internet interconnections and highlighted some possible anti-competitive behaviors of certain ISPs in this arena. Section II covered traditional telephony interconnection regulations and settlements, and showed why such settlement mechanisms are difficult to define and implement for Internet interconnections. Section III will explore the history of FCC regulations concerning information services and cover topics that, while not directly related to Internet interconnections, show that the Internet is forcing the FCC to re-evaluate many of its current regulations relating to information services. The similarity of these issues with respect to ISP interconnections shows that the government has an interest in the continued growth and proliferation of the Internet, which could be hurt by anti-competitive actions regarding such interconnections.

History of Information Services Regulation

Over the years, the FCC has generally taken a "no regulation" stance with regards to data services over traditional telecommunications infrastructures. This section gives a brief review of the FCC decisions relating to information services as described by Computer Inquiries I, II, and III below. These inquiries were FCC investigations of and decisions on how computer related services over the telephone network should be handled.

Computer Inquiry I: Data Service Not Regulated

In the Computer Inquiry I proceeding, which opened in 1966, the FCC requested comments on the "interdependence of computers and communications services.[120]" The Commission issued a final decision in March of 1971, which addressed two important questions: (1) are computer data processing service providers subject to common carrier regulation; and (2) should common carriers that provide data processing services themselves be subject to regulations that prevent them from discriminating against competing data service providers?

For the first question, the FCC answered no, that they should not regulate computer data processing service providers. The decision noted that while the 1934 Communications Act gave them the authority to regulate any services offered by wire or radio, the Act also specified that the FCC remove barriers to market entry when such barriers unlawfully restricted free enterprise. Therefore, "data processing activities" were not regulated from the very early days of the industry, and were therefore exempt from common carrier requirements of the 1934 Act.

For the second question, there were serious competitive concerns if common carriers entered into competition with data service providers. The common carriers could unfairly compete with the data providers because the data providers required the facilities of common carrier in order to be able to operate. If the common carriers themselves provided such services, they would have the ability (and incentive) to discriminate against the data provider in terms of pricing, provisioning, etc. of the necessary facilities. Also, the common carrier could improperly cross-subsidize their unregulated data offerings with revenues from regulated basic services. Therefore, the FCC ruled that common carriers seeking to offer data services could do so only through a separate affiliate. In this way, the FCC required the common carrier to serve their data service affiliate and competitors equally in all respects, so that the affiliate would note receive preferential provisioning, quality, maintenance, or costs.

---

[120] Oxman, Jason. "The FCC and the Unregulation of the Internet." Office of Plans and Policy (OPP), Federal Communications Commission. July, 1999. P. 8.

Computer Inquiry II: Basic vs. Enhanced Service Providers:

With this proceeding, which opened in the 1970's, the FCC focused on the need for a "workable definition of both regulated telecommunications services and unregulated data services.[121]" The FCC issued a final ruling 1980[122], which resulted in the creation of two categories of service providers, "basic service providers" and "enhanced service providers." The term "basic service" was defined as a common carrier offering of a pure "transmission capacity for the movement of information." Basic service is simply basic telephone services. The term "enhanced service" was defined as "services offered over common carrier transmission facilities used in interstate communications, which employ computer processing applications that act on the format, content, code, protocol, or similar aspects of the subscriber's transmitted information; provide the subscriber additional, different, or restructured information, or involve subscriber interaction with stored information."

Present day enhanced services include such things as voice messaging, alarm monitoring, and Internet access services. Enhanced services can be thought of as basic service plus something. For example, a voice mailbox provided by the telephone company requires the basic service of the phone network in order to be able to get to the enhanced service of the voice messaging system. Also, Internet access via analogue dial requires the basic service of a phone company plus the added services of a local ISP.

---

[121] Ibid., p. 10.

[122] Blumenfield and Cohen Law and Technology Group. "Telecom Industry History." Available at URL: http://www.technologylaw.com/telephony.html.

Computer Inquiry III: Removal of "separate affiliation" rule:

In Computer Inquiry III, initially released by the FCC in 1986[123], replaced the "separate affiliate" requirement of common carriers that wanted to provide data services with a model of safeguards. This allowed common carriers to offer enhanced services that could be integrated with basic services. Common carriers were required to unbundle the enhanced service offerings from their basic service offerings. Note that enhanced services were still not regulated in anyway. This decision also reiterated that common carriers had to make basic services available to enhanced service providers at tariffed rates.

Current Regulatory Issues the Internet has Created

As seen in Computer Inquiries I, II, and III, the FCC has basically tried to not regulate any data communications over the past 30 years. In fact, in a recent interview, current FCC Chairman William Kennard stated "Our deregulatory posture has been in place for a long time and will stay in place… The amazing growth of the Internet is testament to the wisdom of our deregulatory approach.[124]" However, former FCC commissioner Glen Robinson has pointed out that competition, which the current FCC commission is trying to produce via the Telecommunications Act of 1996, historically has led to more regulation, not less.[125] The reason for this is that new entrants want a particular type of deregulation that favors them, while the incumbents always feel threatened by new technologies and claim that regulation is needed for a level playing field.

[123] Blumenfield and Cohen Law and Technology Group. "Telecom Industry History." Available at URL: http://www.technologylaw.com/telephony.html.

[124] "Interview: FCC Chairman William Kennard." Internet Industry Directory. URL: http://www.internetindustry.com/Interviews/kennard.shtml.

[125] EMAIL: "Notes from a Symposium of Executives/Regulators." Symposium was called "Changing Times, Changing Technology," and was held on April 17, 1999. Brought together by the Harvard Journal of Law and Technology.

This section will explore some of the regulatory issues the Internet has created that have received FCC attention recently. Most of these issues are not directly related to Internet interconnections – i.e. ISP-to-ISP connections used for the exchange of routing information and traffic. However, they are related in some way to Internet connectivity or Internet effects on traditional telecommunication regulations, and it is useful to understand how the FCC has been dealing with such issues to see how they may deal with ISP-to-ISP interconnections in the future.

Reciprocal Compensation and Access Charges

Recall from Section II that reciprocal compensation is a type of bilateral settlement used between two telecommunications providers for local telephone calls. Local telephone companies are compensated for completing calls that are placed by one of its competitor's customers, and the companies involved are typically the incumbent local exchange carrier (ILEC) (which would have a large customer base) and a competitive local exchange carrier (CLEC) that has recently entered the market (and would therefore have a smaller customer base).

Reciprocal compensation had its start in the Telecommunications Act of 1996, which had many goals, one of which was to introduce competition in the local telephone service provider market.[126] In order to do that, a settlement mechanism for when one local network hands off traffic to another had to be created, and that mechanism was reciprocal compensation. As the network that originates the local call pays the network that terminates the call, money follows the traffic in the sense that the originating network receives money from their customer, and passes some of that money on to the terminating network to pay for that termination.

---

[126] Cannon, Robert. "Legal Developments: FCC release Reciprocal Compensation Order." Boardwatch Magazine. May 1999.

What happened because of this, though, was that the new networks (the CLECs) went after customers that receive many calls, such as pizza delivery stores, flower shops, and especially ISP that provide dial access services. Since the network that terminates the calls receives the settlement, the goal is to terminate as many calls as possible, and therefore high call receivers are the desired customers of the local providers. ISPs have modem banks that receive many, many calls. The ILECs originally wanted the terminating network to receive the money, because, as the incumbent, they would have had the much larger customer base and assumed that would mean they would terminate more calls. However, as soon as the CLECs went after the customers that received a large number of calls, the ILECs cried foul, withheld reciprocal compensation payments to the CLECs and brought the issue before their local regulatory agencies to rule on.

Twenty-six state regulatory commissions (all that reviewed the issue) declared that a telephone company that originates a call to an ISP must pay the telephone company that provides the local service to the ISP reciprocal compensation.[127] Many of the incumbents argued that Internet traffic is not "local" in nature, but should instead be considered "inter-state" because traffic often begins in one state and ends in another state. Such inter-state traffic is not subject to reciprocal compensation per FCC regulations, as it is more like long distance traffic. Long distance companies pay access charges to local phone companies, and when two local phone companies are involved in carrying the call to the long distance provider, the two local companies share the access charges paid by the long distance company and no reciprocal compensation is due. Unlike long distance carriers, ISPs do not pay access charges to local telephone companies because ISPs are considered "enhanced service providers" (ESPs), per the FCC's decision in Computer Inquiry II.[128] In that decision, the FCC concluded that ESPs were not subject to common carrier regulation and

---

[127] FCC Reciprocal Compensation Fact Sheet. February 1999. URL: http://www.fcc.gov/Bureaus/Common_Carrier/Factsheets/nominute.html.

[128] Oxman, Jason. "The FCC and the Unregulation of the Internet." Office of Plans and Policy (OPP), Federal Communications Commission. July, 1999. P.16.

that they did not use the local exchange network in a manner analogous to inter-exchange carriers (ICX's) (i.e., long distance providers).

The issue of whether or not Internet calls should be considered long distance (and therefore subject to access charges) or local (and therefore subject to reciprocal compensation) was brought forward to the FCC in two separate proceedings. First, for access charges, in CC Docket 96-262[129] issued in December of 1996, the FCC sought comments on the issue, and in May of 1997 issued FCC 97-158[130], the Access Reform Order, in which they concluded that the existing rate structure for ISPs should remain in place (i.e. they were not required to pay interstate access charges.) Second, the issue of whether the ILECs should pay reciprocal compensation to CLECs that sold service to ISPs came up in February of 1999, Docket 99-38[131], and the FCC concluded the carriers were bound by their existing interconnection agreements and were therefore subject to reciprocal compensation obligations to the extent of such agreements or as determined by state commissions. The FCC declared that Internet traffic is jurisdictionally mixed in that it appears to be mostly interstate in nature (in which case the FCC has authority to make regulations over it), but that it was still exempt from interstate access charges.

Robert Cannon further broke this down in a Boardwatch article as follows:[132] The FCC analyzed Internet calls in an end-to-end fashion, and determined that no matter what happened in the middle, the end points of the call were all that mattered. In other words, an end users starts a call and terminates it at the Internet destination the user is accessing, and the elements in the middle, such as the ISP, router, upstream providers, DNS servers, etc., are incidental to the complete transmission. Such communications were deemed to be a mix of inter- and intra- state

[129] FCC Access Charge Reform Homepage. URL: http://www.fcc.gov/isp.html

[130] FCC 97-158. URL: http://www.fcc.gov/ccb/access/fcc97158.html

[131] FCC Docket CC 99-38. URL: http://www.fcc.gov/Bureaus/Common_Carrier/Orders/1999/fcc99038.pdf

communications, for which the FCC has no rules to deal with. In the absence of any rule, the FCC declared that "the local telephone companies had 'voluntarily' included this type of traffic within their interconnection and reciprocal compensation agreements," and therefore the FCC would not interfere.

It should be noted at this point that this whole issue was the heart of the now infamous "modem tax" rumor, in that if the ILECs were allowed to treat the ISPs as IXCs and therefore be able to charge per minute access-charges, the ISPs would have to in turn disband the common flat-rate dial pricing scheme generally offered to Internet end users, and instead charge per-minute fees to connect. However, as can be see from the previous discussion, that was never the real reason this was brought forward to the FCC. It was instead an attempt by the ILECs to stop paying reciprocal compensation to CLECs for Internet calls.

Twenty-six states had already decided that reciprocal compensation was due, and the FCC, by declaring Internet calls as "jurisdictionally mixed," sustained states' decisions. So Internet calls are considered local in the sense that reciprocal compensation is required, but are considered inter-state yet fall under the enhanced service provider umbrella such that access charges are not required. This is a prime example where historical and/or current regulations and definitions do not fit the rapidly changing environment of telecommunications, and therefore rule-making decisions on Internet issues is very difficult. As such, the FCC has initiated a rule making proceeding to determine how ISP bound traffic will be handled.[133]

As part of the request for comments on a rule-making procedure, the commission also tentatively concluded that commercial negotiations are the ideal means for establishing the terms of

---

[132] Cannon, Robert. "Legal Developments: FCC releases Reciprocal Compensation Order." Cannon. Boardwatch Magazine. May 1999.

[133] FCC Docket CC 99-2. URL: http://www.fcc.gov/Bureaus/Common_Carrier/News_Releases/1999/nrcc9014.txt.

interconnection contracts (reciprocal compensation in particular) with two proposals.[134] The first proposal is that "commercial negotiations are the ideal means of establishing the terms of interconnection agreements." The second proposal is that the "FCC would adapt a set of federal rules governing inter-carrier compensation for ISP bound traffic, pursuant to which both parties would engage in negotiations concerning rates, terms, and conditions applicable to delivery of interstate ISP-bound traffic." The main difference between the two proposals is that in the first, the states are the arbiter, while in the second, the FCC can become the arbiter based on their predetermined rules. Under either proposal, the negotiation takes place between the two parties, and if there is any dispute, it can be resolved by either the state regulatory agency (in proposal one), or the FCC (in proposal two).

The first proposal is basically the status quo between ILECs and CLECs today. Note that it is also the status quo between ISPs that interconnect – most "peers" negotiate such that neither ISP pays the other, though that is not always the case, and in the case of "transit" the two ISPs negotiate a rate based on a customer-provider relationship. In the second proposal, however, the FCC is proposing to establish a set of rules that "govern" how "ISP" traffic is handled in terms of the ILEC to CLEC hand-off. In this case, the rules are for the local loop to the CLEC that provides the ISP local service, so it is different than ISP to ISP interconnections in terms of "peering," but the jump from one interconnection type to the other may not be that difficult for the FCC to make. Leaving the establishment of inter-carrier compensation frameworks to private contractual negotiations between RBOCs and their competitors has failed miserably,[135] as seen in the preceding discussion of regulatory proceedings. The same can also be said about ISP-to-ISP interconnections, and it may not be long before they too are brought before the FCC as regulatory proceedings. As the FCC has

[134] Ibid.

[135] Geist J., Rudolf. "Policy Forum: Regulation of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

proposed either state regulatory agency or the FCC itself become the final arbiter in ILEC to CLEC issues, it may at some point do the same in ISP-to-ISP issues.

Section 271 and 706 of the Telecommunications Act

RBOCs are currently prohibited by section 271 (b) of the Telecommunications Act of 1996 from providing services within their own regions that extend beyond their local access transport areas (LATAs)[136]. This prohibition was established when AT&T was dissolved into separate companies in the early 1980's, and was established to prevent the RBOCs from providing long distance services. The RBOCs had been created specifically to provide local services, leaving AT&T and its competitors to provide long distance services. However, in light of the current Act's goal of more competition through deregulation, the RBOCs are now permitted to provide long distance services and any other services that extend beyond their own LATA's providing they met certain prescribed criteria. The Act includes a 14-point "competitive entry" checklist that the RBOCs must pass before the FCC will grant approval for such services which includes such things as the following:

- interconnection with local competitors as defined by section 251(c)(1) and 252(d)(2) (i)
- non-discriminatory access to poles, ducts, conduits, and rights of way for local competitors (iii)
- unbundling of the local loop (iv)
- reciprocal compensation agreements as required by section 252(d)(2) (xiii)

(The Roman numeral in parentheses following the list entry is the number in the Act's 14-point checklist.) The list goes on, but these four items are the main clauses relating to interconnections

[136] http://www.fcc.gov/Reports/tcom1996.pdf. p. 39.

and settlements. Although not necessarily the same kind of interconnection and settlement issues involved in ISP-to-ISP communications, these do indicate that the FCC has an interest in regulating interconnections and settlement mechanisms between competitors in the local market.

The other aspect of this regulation that is more directly related to Internet interconnections is the fact that RBOCs cannot provide data services (e.g. Internet services) between LATAs, which basically translates to "RBOCs can not develop national backbone infrastructures for providing Internet services." Instead, the RBOCs are forced to buy transit connections from national providers in order to get full Internet connectivity to provide end-to-end service for their customers.

The RBOCs have tried various avenues to get around this checklist in order to provider inter-LATA services, including data services, such that they could build national backbones. In "The Need for Facilities-Based Internet Backbone Competition,[137]" which was authored by Robert Gibson but sponsored by the RBOCs under petition of section 706 of the Telecommunications ACT, "seeks regulatory authority from the FCC to deploy regional Internet backbones, i.e. broadband inter-LATA data services." Section 706 of the Act states "The Commission and each State commission with regulatory jurisdiction over telecommunications services shall encourage the deployment on a reasonable and timely basis of advanced telecommunications capability to all Americans…[138]" Gibson covers much of the same ground that this thesis has covered to this point – namely the potentially anti-competitive behaviors, concentration of control of facilities such as NAPs, and the fast consolidation of ISPs through corporate buyouts and mergers – in addition to section 706, in order to petition the FCC to override the restriction placed on RBOCs of offering inter-LATA services. Gibson's main point is that by allowing the RBOCs to offer inter-LATA data services, which in effect would allow traffic to flow between an RBOC's local service areas without

[137] Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

[138] The Telecommunications Act of 1996. http://www.fcc.gov/Reports/tcom1996.pdf. p. 120.

first having to pass through a national ISP, would provide faster local services and more competition in the Internet access market.

Such complaints have not fallen on deaf ears within Congress. In the "Internet Freedom Act of 1999,[139]" Senators Boucher and Goodlatte have proposed lifting such inter-LATA restrictions. This proposed Act states "the FCC is currently ignoring its responsibility under the Telecommunications Act of 1996 to provide regulatory relief to incumbent phone companies [RBOCs] by removing existing regulations on data traffic. Existing suppliers of Internet backbones are simply unable to keep up with the demand for high-speed, high capacity backbone bandwidth." Section 202 of The Internet Freedom Act of 1999 would amend the Communications Act such that the restrictions of section 271 would only apply to voice traffic, thereby opening the market for inter-LATA data service for RBOCs.[140]

Note again the dichotomy between basic services (voice communications) and enhanced services (services beyond voice which include data services). In this case, the Act of 1996 has ignored the dichotomy in basic and enhanced services by restricting the RBOCs from providing inter-LATA services for both voice *and* data, whereas the proposed "Internet Freedom Act of 1999" would impose the dichotomy so that data services would be freed from such regulation. It is clear that the dichotomy is causing more confusion in the current telecommunications market of rapid change than it is relieving and that new definitions are needed.

Also note that RBOCs are currently forced into buying transit services from national backbone providers. By lifting this regulation and allowing RBOCs to build national their own Internet backbone infrastructures, the RBOCs could potentially negotiate settlement free peering

---

[139] Cannon, Robert. "Legal Developments: Congress Addresses Internet Access, Spam, and Cyberporn." Boardwatch Magazine. July 1999.

[140] Section-by-Section Summary of the Internet Freedom Act of 1999. URL: http://www.house.gov/goodlatte/sectsum.htm

with other national providers, and therefore add more competition to the market. Again, these proposed changes in current regulations shows that ISP-to-ISP interconnections are at the fore of government concerns, which may mean such connections are regulated in the future.

Common Carrier Status

Soon after the first announcements by large ISPs such as UUNET and Sprint that they would terminate their "free peering" agreements with smaller networks (see Section I, "The End of Open Peering") many people began suggesting that ISPs should be required to interconnect to one another. Common carriers (traditional telephony service providers) are required to interconnect with each other, per section 251 (a) of the Communications Act of 1934. However, ISPs are not required to interconnect with each other, which stems semantically from the fact that they are not considered common carriers[141].

Recall from section I, that David Holub lost his job as vice president of Whole Earth Networks over the dispute caused by UUNET's announcement of their intent to end free peering relationships with smaller networks. Soon after, Holub posted his arguments as to why ISPs should be regulated by the government to interconnect. These posting were the among the first arguments that state ISPs should apply for common carrier status so that interconnections would be required by government regulation. Holub's argument follows the following reasoning:[142] Universal reachability is a fundamental tenet for all public networks. The assumption is that a customer who pays for a connection to a public network will receive connectivity to all points on that network.

[141] Frieden, Rob. "Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization." Virginia Journal of Law and Technology. Fall 1998.

[142] Cook, Gordon. "David Holub Explains Model of an Open Peering, Open Interconnect, 'Bill and Keep' Internet." The Cook Report on the Internet. July, 1997.

This is true and has been required of the traditional telephony network, and it is assumed to be true of the Internet. As these Internet is not a single entity, but instead comprised of many networks that connect together, there must be interconnection requirements in order to keep the assumption of universal reachability in tact.

Holub went on to argue that ISP-to-ISP interconnections should be classified as basic telecommunications rather than enhanced services, such that Internet interconnections would be required under section 251 (a) of the Communications Act. To this end, he went so far as lobbying for common carrier status for his ISP Whole Earth Networks (before he was fired) in order to ensure interconnections with other providers. In this regard, Holub was a pioneer on this front as this argument is still used today. However, the FCC has yet to decide that Internet traffic is basic a basic telecommunication service, so it still lives under the umbrella of enhanced services and is therefore not subject to regulation. Therefore, ISP interconnections are not a requirement, and to date have only been held together by market pressure. However, if the tenet that universal reachability is fundamental to the Internet were true, then it would seem that ISP-to-ISP interconnections should be enforced by something stronger than just market demands.

"Last Mile" Open Access: RBOCs and Cable (or xDSL and Cable)

The "last mile" – the wire that connects individual residences to either the local telephone or cable network, has been a very hot issue over the past few years. There is currently a strong push by service providers to install broadband bandwidth to the home because the demands of using the Internet effectively have continued to rise. In the past, the last mile has typically been provided by a monopolistic company[143] – the RBOC for the copper wire used for local telephone service and the

[143] Oram, Andy. "ISPs and Internet Policy: The New Agenda." The American Reporter, Vol. 5, No 1037. March 30, 1999. URL http://www.american-reporter.com/1037/1.html.

cable company for cable television service. RBOCs are scrambling to provide broadband bandwidth over Digital Subscriber Loops (DSL) while cable providers have been just as busy trying to provide broadband bandwidth over cable modems. The RBOCs have been traditionally forced by federal regulations to provide "open access" to their facilities to CLECs[144], whereas cable companies have not been subject to the same open access requirements. This section will examine the regulatory issues the Internet brought forth in this area.

The Telecommunications Act of 1996 had a primary goal of increasing competition such that consumers would benefit by lower prices and increased access to new services. It could indirectly accomplish this by (ironically) decreasing the number of communications companies by allowing mergers of large companies with traditionally different services (i.e. AT&T, a telephone company, buying TCI, a cable provider).[145] The Act deliberately relaxed limits on multiple ownerships (restrictions on how many different services a single company could provide) and entry into various markets so that such competition creating consolidation could occur.

Until recently, RBOCs were required by the Communications Act of 1996 to unbundle network elements to CLECs, such that the CLECs could effectively sell local services at rates that allowed them to be competitive with the RBOCs. However, on September 15th, 1999, the FCC issued an order[146] that revised the Act's original order. Much of this new order dealt with traditional telephony related issues, for which six of the original seven network elements must still be unbundled, but one part of it dealt directly with broadband access to the Internet. The Commission declared that, except in limited situations, the ILECs were not required to unbundle the facilities used to provide high-speed Internet access and other data services – specifically noting

[144] Section 251 (a) of the Communications Act of 1934.

[145] Oram, Andy. "ISPs and Internet Policy: The New Agenda." The American Reporter, Vol. 5, No 1037. March 30, 1999. URL http://www.american-reporter.com/1037/1.html.

[146] "FCC Promotes Local Telecommunications Competition: Adopts Rules on Unbundling of Network Elements." September 15, 1999. URL: http://www.fcc.gov/Bureaus/Common_Carrier/News_Releases/1999/nrcc9066.html

packet switches and digital subscriber line access multiplexers (DSLAMs). The order stated "given the nascent nature of this market and the desire of the Commission to do nothing to discourage the rapid deployment of advanced services, the Commission declined to impose an obligation on incumbents to provide unbundled access.[147]"

This order effectively put RBOCs and cable companies on equal ground. Until this order, RBOCs were required to unbundle components used for high speed Internet access, whereas cable companies that were offering high-speed access over cable modems were not. This had been seen as a situation that was fostering anti-competitive behaviors by cable companies, in that they could provide high speed Internet access over cable but other ISPs in the same geographic area could not. Before this order was issued, the ILECs had been making it extremely difficult for CLECs to gain access to such services by "dragging their feet" in following the regulations imposed by the Act and fighting the open access requirement for high speed Internet access[148].

In addition, cable companies had refused to open up their facilities to competitors on the grounds that they had paid for the upgrades necessary to the cable plant to provide such services, and that they would discontinue the upgrades if they were forced to share the facilities. John Petrillo, executive vice president of AT&T summed this sentiment up: "don't make us open our cable network to competing ISPs. We'll just stop offering our data services instead, because we can't justify the billions of dollars it would require to upgrade the cable network if we had to share it with others. After we've built the infrastructure and we're up and running, society can then decide whether it would be best to open our network to others – but not now.[149]" However, competitors claimed that by the that time, the cable companies would have had an unfair advantage

[147] Id.

[148] Oram, Andy. "ISPs and Internet Policy: The New Agenda." The American Reporter, Vol. 5, No 1037. March 30, 1999. URL http://www.american-reporter.com/1037/1.html.

[149] "Notes from a symposium of executives/regulators" Symposium was called "Changing Times, Changing Technology," and was held on April 17, 1999 and was brought together by the Harvard Journal of Law and Technology. Posted to Cyber-telecom mailing list.

in building a huge customer base, and that those customers would not likely quickly switch over to the new competition once they were allowed to share the infrastructure.

The OpenNET Coalition is a group of ISPs and other interested parties that have banded together to fight for the unbundling of cable companies high-speed data networks, which would allow ISPs other than the cable companies to provide Internet services over the cable plant. They are now likely to fight the FCC's recently declared order to not force ILECs to unbundle similar high-speed data network equipment. When AT&T purchased TCI, each local community had to approve the local cable license transfer. Several communities, including Fremont, CA, Broward County, FL, and Portland, OR, declared that AT&T would have to open their cable plant to competitors. The FCC has taken the stand that such community based regulations may have good intentions but are bad regulations: "Establishing a world where each local community can set the technical standards … dooms the Internets lofty goals to failure.[150]" This follows the FCC's sentiments from two separate proceedings earlier in 1999, in which the FCC stated "The market, not the government, should bring broadband to all Americans.[151]"

As can be seen by the proceeding discussion, it is apparent that the FCC is trying to maintain their position of the "unregulation" of the Internet, by not forcing ILECs and cable companies to open their facilities to competitors. While this seems to go against their stated goals of fostering competition, in some senses it is not. As both ILECs and cable companies dragged their feet in opening access to others, the deployment rate of high-speed advanced services was slow. The FCC has taken the stance that there is competition between the RBOCs and the cable companies giving consumers a choice for broadband providers and has therefore not declared that both must open access to non-facilities based ISPs.

[150] McCarthy, Bill. "Editor's Notes." Boardwatch Magazine. June 1999.

[151] "Legal Developments: FCC Finds Broadband Deployment to be 'Reasonable and Timely.'" Robert Cannon. Boardwatch Magazine. April 1999.

In what can be considered an attempted reversal of the government's stance, Senator Bob Goodlatte has introduced the Internet Freedom Act of 1999.[152]

Section 101 of this proposed Act would:

> "ensure competition among providers of high speed Internet service by requiring that incumbent local telephone companies provide conditioned unbundled local loops as provided by the Telecommunications Act of 1996 when economically reasonable and technically feasible."

In addition, section 102 states the following:

> "would ensure that other broadband access transport providers, in addition to incumbent local telephone companies, allow broadband Internet service providers to compete fairly over their facilities."

Finally, section 103 states the following:

> "it shall be unlawful for a broadband Internet transport provider to engage in unfair methods of competition or unfair deceptive trade practices, the purpose or effect of which is to discriminate in favor of a service provider associated with that transport provider or restrain unreasonably the ability of a service provider not affiliated with a broadband Internet access transport provider to compete in its provision of Internet services."

Whereas the FCC has stepped back on the unbundling of high-speed local access equipment, the Internet Freedom Act is actually pushing forward the regulation of such unbundling requirements. It appears that the FCC is trying to speed the deployment of broadband Internet, by not forcing the owners of such facilities to share those facilities, as those owners have stated they

[152] Section-by-Section summary of the Internet Freedom Act of 1999. URL: http://www.house.gov/goodlatte/sectsum.htm

would simply not build it if they were required to sink all the money into the infrastructure just to have to share it. On the other hand, Congress is trying to foster competition in broadband Internet access by forcing the owners/builders of the infrastructure to share it with others.

For both the FCC and Congress, it is clear that the U.S. government has an interest in both the deployment of the Internet and competition in providing access to the Internet. In this "last mile" issue, there is a fight to ensure broadband access is built and ensure that there is competition without anti-competitive behaviors. These are the same issues that ISP-to-ISP interconnections face today – namely apparent anti-competitive behaviors, the need to deploy the services at fast rates, and true competition. It is clear that the FCC and Congress are very interested in such areas and are watching them closely to ensure true competition and fast build out. Therefore, it is likely that if any behavior in the area of ISP-to-ISP interconnections goes against this grain, the FCC would likely intervene.

Convergence and its Effects on "Enhanced Services" and the Universal Service Fund

In the past several years, the telecommunications industry has seen a tremendous consolidation in terms of company acquisitions and mergers in which the two parties have traditionally offered vastly different services. The AT&T acquisition of TCI is just one such example, in which a traditional telephony service provider purchased a cable television service provider. The reason for this is that what were once disparate services are converging onto a single network. AT&T wants to provide local telephone service, cable television service, and broadband Internet access all over a single "wire" into the home. Other examples in which differing telecommunications service companies have merged/been acquired are AOL/ Teleport and US

West/Cablevision (which became Media One).[153] Such mergers have been the result of telecommunications deregulation and are leading to the fact that telecommunications services are becoming a commodity -- as supply outstrips demand, prices plunge, sometimes to the point of flat-rate services.[154] In addition, IP telephony over the Internet is growing at a tremendous rate, as Internet speeds continue to grow and the quality of such connections improves.

The various telecommunications networks, including the telephony network, the cable television network, and the Internet, all have the foundation to merge or become indistinguishable from each other.[155] The Internet has the capacity and versatility to become a one-size-fits-all transmission and information services medium. In other words, as the content delivered over the various networks becomes digital, it is possible for one network to carry differing services while being transparent to the end users and/or their applications. This fact of converging media has several consequences. For one, it makes it increasingly difficult for the FCC to determine where telecommunications services end and information services begin,[156] which is the very heart of the dichotomy the FCC uses to determine whether or not services are regulated. It also may have dramatic affects on the Universal Service Fund, which was briefly described in Section II. Both of these issues will be discussed further here.

The first issue of convergence is the FCC's distinction of telecommunications as "basic service" and of information services as "enhanced service." The FCC has used this dichotomy since Computer Inquiries I, II, and III, in order to not regulate information services. In Computer Inquiry II[157], the FCC defined "basic services" as the transmission capacity for movement of

---

[153] Ellis, John. "The Forecast for Phone Rates: Flat." St. Petersburg Times, February 27, 1999.

[154] Ibid.

[155] Frieden, Rob. "Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization." Virginia Journal of Law and Technology. Fall 1998.

[156] Oxman, Jason. "The FCC and the Unregulation of the Internet." Office of Plans and Policy (OPP), Federal Communications Commission. July, 1999. p.22

[157] Ibid.

information, or basic telephone service. Enhanced services were defined as "services offered over common carrier transmission facilities … which employ computer processing applications that act on the format, content, code, protocol, or similar aspects of the transmitted information." However, with the convergence of the various networks and the types of transmissions they are able to handle, it is becoming increasingly difficult to determine where basic telecommunications services end and enhanced services begin. Because of this difficulty, the FCC will have to come up with new methods for classifying services in order to help them understand whether regulations are needed or not.

The Universal Service Fund, as described in section II, subsidizes telecommunications services for low-income individuals as well as the costs associated with bringing services to high cost rural areas. Currently, all providers of telecommunications services are required to contribute a portion[158] of their revenues to this fund. However, just as ISPs are exempt from access charges (because they have been classified as enhanced service providers), they are also exempt from contributing to the universal service fund. Because of convergence, IP telephony is becoming more and more popular, and as the quality continues to improve, the Internet is likely to carry an increasing amount of the voice conversations traditionally carried on the telephony network. There are clearly financial reasons to move voice calls to the Internet. Because the telephony network is classified as a basic service, it is subject to regulations, including access charges, reciprocal compensation, and universal service fund contributions. Therefore, voice calls on the telephony network have costs associated with them that voice calls over the Internet do not have. Today, Internet calls are cheaper (in part) due to the lack of regulation, so more and more calls will migrate

[158] The USF contribution factor is calculated quarterly by the FCC based on the ratio of total projected quarterly costs of the universal service support mechanisms to total end-user interstate and international telecommunications revenues. The latest rates can be found at the following URL: http://www.fcc.gov/ccb/universal_service/quarter.html.

to the Internet as the quality continues to improve.[159] As this happens, the USF will lose some of its contribution base and may begin to dry up.[160]

In addition, the Communications Act of 1996 expanded the telecommunications universal service mission to include "access to advanced telecommunications and information services ... throughout all regions of the Nation," and further expanded the beneficiaries of the fund to include schools and libraries. The FCC further declared that Internet access was part of the "e-rate" telecommunications discount for schools and libraries, meaning that such institutions would be subsidized so that they could purchase Internet access. Although individual households were not included in Internet access subsidies via the universal service fund, the FCC did not rule this out and it may happen as traditional services migrate to the Internet, which would make universal access to the Internet more important.

In fact, in April of 1998, in response to congressional direction, the FCC released a report on the Commission's progress in implementing the universal service provisions of the 1996 Act[161]. The FCC concluded that information service providers (i.e. ISPs) would not be required to make direct contributions to the universal service fund as telecommunication carriers are required to do. However, they also expressed discomfort with maintaining a blanket exemption on all types of Internet telephony. "Certain phone to phone IP telephony services lack the characteristics that would render them information services within the meaning of the statute, and instead bear the characteristics of 'telecommunications services' as defined by the Act of 1996[162]." However, in this report, the FCC did not take further action. However, their statements do show that the matter is

---

[159] Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998. p.13

[160] Geist J., Rudolf. "Policy Forum: Regulation of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

[161] http://www.fcc.gov/Bureaus/Common_Carrier/Reports/fcc98067.pdf

[162] Ibid.

under consideration. Rudolf Geist, writer of Boardwatch Magazine's Policy Forum, writes that at some point the FCC "may be forced to classify all providers who use the Internet platform as 'telecommunications carriers' regardless of the type of service they provide, to ensure that a significant USF contribution base remains."[163]

To summarize, convergence causes two major issues. First, it will no longer be possible to live with the dichotomy of basic service vs. enhanced service, as it is increasingly difficult to determine where basic telecommunications services end and where enhanced information services begin. Second, as traditionally regulated services migrate to the unregulated Internet, the Universal Service Fund base may dry up to the point where low-income individuals and high-cost rural areas are no longer subsidized enough to maintain telecommunications services. Finally, it must be pointed out that as telecommunications services converge to one network, and with it looking more and more like that one network will be the Internet, guaranteed connectivity from any point on the network to any other point on the network is a necessary. Therefore, regulated interconnections between ISPs become desirable, as forced interconnection may be the only way to guarantee such end-to-end connectivity.

---

[163] Geist J., Rudolf. "Policy Forum: Regulation of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

Section IV: The Future of Internet Interconnections

Thus far, this thesis has shown that what was once an open peering policy in which most networks interconnected with other networks on a settlement-free basis has become more of a hierarchical discriminatory model in which large networks try to force smaller networks into paying for such connections. These negotiations often take place under non-disclosure agreements (NDA's) in which the criteria for settlement free interconnections are not clearly defined by the large ISPs. These large ISPs try to set the terms of the interconnections to what may appear to be outside the realm of normal competitive pricing constraints without the natural oversight that comes with a marketplace where essential price information is known amongst rivals.[164]

The crux of the problem is how to have a competitive Internet where there can be new entrants to the market and where smaller providers are able to innovate and introduce new products, without hurting large networks that have invested large amounts of money in their infrastructures. This is a difficult problem to solve and most of those involved have their own opinions on how best to do it. "It seems that everyone agrees peering needs change, and they unanimously agree that no one is going to agree on anything.[165]" We have seen that the Internet does not easily fall into a model of settlement-based interconnections such as that used in the traditional telephony world, and that determining a fair settlement mechanism is not a simple task or one that is likely to be agreed on by all involved parties.

[164]Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html. p.9

[165] Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes, May 1998. p.11.

The evolution of public networks suggests that monopolistic abuses are commonplace unless kept in place, either by competitive market forces or by some form of regulation.[166] To date, market forces have been able to preserve the basic tenant of end-to-end connectivity of the Internet, but this may not continue as mergers continue and dominant players become more dominant. Universal reachability on the Internet is in the public interest and needs to be guaranteed in some way. This section looks first at the possibility of industry led "self-regulation," and then at the likely-hood of government regulation.

Industry "Self Regulation"

This section will explore several ways the Internet industry could potentially "self-regulate" itself such that government intervention does not become necessary, including open frameworks, multi-lateral peering arrangements (MLPA's), and Mike Gaddis's proposal of the Brokered Private Peering system. All three ideas propose ways in which negotiations for interconnections would be made more objective rather than the current subjective state. Objective determinations would leave no questions as to which ISPs should peer on settlement-free basis with which other ISPs and which should interconnect in a customer-provider relationship.

Open Framework

A major problem with the current state of interconnections is that there is no open framework on peering requirements. Large networks may have once publicized peering criteria, such as 24x7 operator coverage, connections in geographically disperse locations, and a certain size level of the

[166] Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html. p.9

potential peer's backbone, but more often than not, such terms are now only discussed under NDA's. Because of this, it is increasingly difficult for a small network to grow into a national tier one network that is able to get settlement free peering with all of the major networks.

An open framework in which all networks disclosed their peering requirements could allow all networks to have adequate information to use for decision making on peers and to better evaluate their commercial positions relative to the industry.[167] This would also ensure to some degree that commercial discrimination does not take place, as all networks would have the same information in front of them when negotiating interconnection agreements.

Such an open framework could be initiated in the form of an Industry forum, which would be more in spirit with the past evolution of the Internet (more so than government regulation). In addition, the U.S. has tried to follow such a market led approach, and a recent White House report entitled "A Framework for Global Electronic Commerce[168]" seeks a non-regulatory, industry led system to govern the Internet from an international perspective as well (as opposed to government led regulation).

The idea of an open framework could help the situation, but the likelihood of it happening without government intervention is remote. Most large networks are happy with their current positions when it comes to negotiating interconnection agreements. As it stands, with their requirements not publicized and negotiations only occurring under NDA, the requirements can be changed to meet the varying situations that occur with each discussion for every single potential interconnection. In this way, they have an advantage over every possible peer and are they are very unlikely to give that up via a publicized open framework of requirements.

---

167 Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

168 A Framework for Global Economic Commerce. URL: http://www.ecomerce.gov/framework.html.

Multi-lateral Peering Agreements

Another potential way the industry could self-govern itself is through multi-lateral peering agreements (MLPA's). When two networks interconnect at private peering points and/or at most NAPs, they negotiate a bi-lateral peering agreement, which is a private contract between the two parties that defines the terms of the interconnection. As the name implies, a multi-lateral peering agreement is a contract that defines the terms of an interconnection between multiple parties. Such arrangements obviously only make sense at public exchange points where multiple networks can interconnect to each other.

Some NAPs require all networks that want to participate at that NAP to sign an MLPA, which defines the terms of the NAP connection and the interconnections that will be established. In such cases, the MLPA may strictly require all networks to interconnect with all others that are at the NAP, or it may define some criteria that defines which networks are required to interconnect with which other networks. Other NAPs have MLPA's but don't require all networks that participate at the NAP to sign it. In such cases, the MLPA may define that all networks that signed the MLPA must peer with all others, or it may define criteria that defines which networks are required to peer with other networks. In the case of networks that join such a NAP but don't sign an MLPA, those networks are required to negotiate bi-lateral peering agreements (BLPA) with any other network at the NAP with which they want to establish an interconnection.

The first example of an MLPA NAP was the Commercial Internet eXchange (CIX). Early CIX arrangements were based on a description of the infrastructure of each party, in which acknowledgements of peer capability were based on the operation of a national transit infrastructure of a minimum capability, but that was later changed such that payment of the CIX association member fee would give peer status.[169] In effect, each ISP had a single agreement with the CIX association, which in turn meant they had the ability to peer with all other association members.

In almost every case, MLPA's have failed to ensure all networks interconnect to one another. NAPs that have MLPA's often don't enforce them such that networks are free to pursue BLPA's with those other networks that they choose to. For those NAPs that did require MLPA's, certain networks that were not interested in being forced to peer with all others would simply not join. In the case of the CIX, which was the first public exchange point and at one point was quite large, it simply dwindled in size as larger providers let their contracts expire (often without implementing necessary upgrades to prevent latency and loss!) and then pulled out. In short, MLPA's often mean that certain networks will not participate in those NAPs as they lose the ability to privately negotiate interconnection terms on a bi-lateral basis. Therefore, many networks feel that they lose their advantages of being able to negotiate under NDA's.

Brokered Private Peering

Michael Gaddis, former CTO of the ISP Savvis, along with contributions from ISPs Electric Lightwave, Exodus, and Williams, has authored a white paper describing the "Brokered Private Peering Group (BPP)." As defined in their words, the BPP would be "an organization dedicated to building a lasting Internet data inter-exchange architecture built upon a sound commercial foundation that recognizes the importance of high quality, scaleable, peer-to-peer inter-exchange bandwidth between Internet Service Providers."[170] The introduction of the white paper details many of the same problems with peering that have been discussed in this thesis -- namely that the

[169] Huston, Geoff. "Interconnection, Peering and Settlements – Part I." Internet Protocol Journal, Volume2, Number 1, March 1999. Soft copy available at URL http://www.cisco.com/ipj

[170] The white paper can be found at URL http://boardwatch.internet.com/mag/98/may/bwm29.html

public exchange points have struggled to keep up with the growth of the Internet, that no interconnection guidelines exist for which networks should peer with other networks and under what conditions those interconnections should be made, etc. It also explains that the costs of private peering are a hindrance to smaller ISPs or those ISPs that are not facilities-based networks because the costs present a barrier for them to enter into private peering arrangements, and are therefore subject to the "bad bandwidth" at the public NAPs.

Much of the BPP proposal has to do with the formation of a scaleable ATM exchange architecture whereby "three or more peering partners interconnect to an ATM switch in distributed locations."[171] In this architecture, there is a DS3 (minimum), OC3, or OC12 loop running multiple virtual channels, one for each peer. It is a shared interconnect with dedicated bandwidth, at fractional physical circuit speeds that use virtual circuits (VCs) between two peers to act as a dedicated local loop. This architecture gives the benefits of private peering (in this case, namely "good" private line-like bandwidth) without local loop costs for each ISP at the exchange because all connected networks share one large local loop, with virtual connections logically giving the appearance of multiple direct connections. Within this framework, there are rules governing utilization such that when a VC between two peers reaches certain levels, it must be upgraded. This framework is administered by the BPP group, thereby ensuring a certain level of quality throughout the exchange. The goal of the architecture is zero packet loss, low delay, and high availability.

The architecture of the BPP exchange does help with certain issues of peering, such as the quality of the NAPs, the responsibility of the ISPs to upgrade as needed to ensure quality, etc., but it goes further than that. As we have seen, one of the main problems with interconnections is that there are no objective criteria to determine which networks should peer with other networks, but instead is left to the subjective views of each provider. The BPP plan goes beyond a simple

[171] Ibid, p.3.

architecture and attempts to objectively define which networks must peer with each other. It accomplishes this with a primary and secondary classification system, as described below, that attempts to define "true peer" status.

The primary classification system defines three types of ISPs based on geographic coverage: national, regional, and local. A national provider must have a national network infrastructure (owned or leased) with the ability to provide commercial Internet access or transports to customers coast to coast within the US, have presence and be able to sell commercial access in four of the six BPP defined regions, have a professionally staffed Network Operations Center (NOC) that is manned 24 hours a day 365 days a year, and have the engineering skills to operate a distributed peer interchange with multiple providers. Regional providers are defined as not meeting the requirements of a national provider, yet have presence in one or more of the BPP defined regions. Local ISPs are those providers that do not meet the requirements of the national or regional definitions.

The secondary classification places all BPP members into one of the following groups: business access oriented, consumer dial-in access oriented, or web centric. Service providers are defined as web centric if they send more data than they receive by a factor of three to one through the BPP exchange point. Service providers are considered as business access oriented if it is not web-centric and 25% of its gross revenue is from businesses or other ISPs. And finally, service providers are considered dial-centric if they primarily sell (i.e. greater than 50% of their gross revenue) dial-in, cable modem, and/or DSL type Internet services to consumer oriented customers. (Note that a service provider could be classified as both business access and consumer dial-in access oriented under these definitions).

The BPP then goes on to defined minimum requirements for true peer status. When two ISP's are classified as true peers, they are required by the BPP to establish interconnections. There are three true peer groups defined, national access class, national dial-in class, and national web-centric class. The requirements of these groups are: national access class must have a minimum of 1000 revenue generating business access circuits at fractional T1 speeds (minimum 256Kbps) or above, and have presence in at least one BPP exchange in all six of the BPP regions; dial-in class must have a minimum of 300,000 revenue generating consumer dial-in (or equivalent) accounts; web-centric must have a minimum of 200 dedicated computer web-centric host accounts that collectively push greater than 500Mbps for more than five hours per day. When an ISP reaches true peer status, it must establish peering relationships with other providers. Smaller ISPs who do not meet the requirements of true peer status have the option of trying to negotiate free peering (outside of the MLPA of the BPP), paying for peering, or buying transiting connections from other networks.

Finally, the BPP white paper defines an "inter-class" true peering classification system. Here, an ISP of "national access true peer class" that has the ability to peer in all six BPP peering regions is defined to be a true peer to national dial-in and national web-centric classed providers. Similar definitions are given for national dial-in true peers (must peer with true peers of national access class and national web-centric class ISPs if they have the ability to peer in all six BPP regions) and national web-centric true peers (must peer with national access class and national dial-in class true peers if they have the ability to peer in all 6 BPP regions).

The main point of all of these classifications and mandated peering requirements is an attempt to objectively define which providers must establish settlement free interconnections with other providers. There is always room for voluntary interconnections, whether on a settlement free arrangement or on a customer/provider relationship. This overcomes one of the main problems with Interprovider interconnections as they stand today, in that each provider today subjectively decides whether to interconnect with another provider on a network-by-network basis, and therefore such interconnections follow no deterministic model.

The BPP initiative has currently stalled, as large networks apparently see nothing to gain from joining it.[172] In fact, they probably saw a lot to lose. In the current state of affairs, large networks have the advantage of negotiating interconnection agreements under NDA's in which, as has been discussed, they often have a decided advantage. By moving to the BPP, they would be required to peer with other providers in which they may have otherwise negotiated some kind of provider relationship in which they would have been by the other provider for the interconnection.

Why Industry Self-regulation Will Fail

Open frameworks, MLPA's, and initiatives such as Brokered Private Peering have all failed, and there is not likely to be any industry led proposal that will succeed in enforcing Internet interconnections. The reason for this is that large networks are currently able to negotiate payments for interconnections with other networks, and they could potentially lose this revenue stream if settlement free interconnections of any kind were required. In all cases of self-regulation, networks would have to share more information than they do today, thus giving away what they see as an advantage in any negotiations. The industry simply has not been able to regulate itself with respect to fair and non-discriminatory mechanisms for the exchange of traffic[173] and that is not likely to change.

[172] Angel, Jonathan. "Toll Lanes on the Information Superhighway." Network Magazine incorporating Data Communications. February 2000.

[173] Geist J., Rudolf. "Policy Forum: Regulation of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

Conclusions: Government Regulation

To date, the FCC has maintained the separation of "basic" traditional telephony services and "enhanced" information services, which means that information services have remained unregulated. Therefore, interconnections between Internet Service Providers have not been and are currently not regulated. However, as admitted by the FCC, "the distinction between the traditionally regulated and unregulated sectors of the industry is blurring.[174]" In a recent Office of Plans and Policy (OPP) paper entitled "The FCC and the Unregulation of the Internet,[175]" Jason Oxman documented the history of how the FCC has tried to not regulate information services and instead let a non-regulated, market driven industry evolve. The paper's recommendations to the FCC are to maintain such a stance, but do leave open the door for possible regulation if necessary:

1. Do not automatically impose legacy regulations on new technologies.
2. When Internet-based services replace traditional legacy services, begin to deregulate the old instead of regulate the new.
3. Maintain a watchful eye to ensure that anti-competitive behavior does not develop, do not regulate on the perception of potential future bottlenecks, and be careful that any regulatory responses are the minimum necessary and outweigh the costs of regulation.

The key point is that the FCC must ensure that anti-competitive behaviors do not develop and that if they do, any regulations they develop are the minimum necessary. The question is whether or not any of the industry's actions to date can be classified as anti-competitive. In light of the fact that so far market pressures have kept the fundamental tenant of end-to-end connectivity in tact means that the actions which have occurred under NDA's have not yet had to be reviewed by the Justice

[174] Oxman, Jason. "The FCC and the Unregulation of the Internet." Office of Plans and Policy (OPP), Federal Communications Commission. July, 1999.

[175] Ibid.

Department or the FCC to determine if such actions are anti-competitive. However, peering policies and peering negotiations of ISPs will likely be reviewed at some point because so many small ISPs have complained about possible anti-competitive behaviors of large ISPs. Once such a review occurs, if the Justice Department and FCC find that market entry or growth have been hindered by the actions of certain large ISPs, they are likely to regulate Internet interconnections.

It is clear that the FCC will become involved with Internet issues to maintain a competitive market place as is evidenced by several of their actions to date. In the case of how the Internet is affecting access charges and reciprocal compensation of traditional telephony providers, the FCC issued a ruling that Internet traffic is jurisdictionally mixed in nature, meaning that it appears to be mostly interstate, in which case the FCC would have the authority to regulate it, but that it was still exempt from access charges. Basically, the FCC had no current classifications for which such issues fit into. The same is true of many services that are offered over the Internet today or will be offered over it shortly, such as high quality voice over IP. In such cases of convergence, the FCC cannot easily classify such services as "basic" or "enhanced," and can therefore not fit them into their current regulated or non-regulated boxes.

As telephony services migrate to the Internet, several factors will make the FCC seriously consider regulation of interconnections. Universal Service Fund contributions are only required of providers that offer "basic" services. Those individuals that have ISPs that offer voice services or those that use software that provides voice over IP are likely to discontinue traditional local telephone services as the quality of voice over the Internet improves. Therefore, over time, the Universal Service Fund may deplete to the point that offering rural customers phone services for prices similar to those in urban areas is no longer possible. Offering such service to all citizens has been a basic tenant of the government because lack of such service is seen as a social and economical barrier between those who can't afford the service and those who can.

In addition, affordable local phone service (made possible for some only via the Universal Service Fund) provides universal reachability to everyone else that has phone service. Such reachability is vitally important for such things as emergency services via 911. In the case of the Internet, universal end-to-end connectivity is only occurring due to market pressures. Having a vital service kept in check by something as unstable as market pressure is clearly not a good idea. Therefore, as a critical mass of users migrate basic phone service to the Internet, ensuring universal reachability will be deemed necessary by the FCC, and this can only be done via regulations of interconnections.

It is certainly interesting to recall the case of how the Internet is affecting access charges and reciprocal compensation. The FCC has stated that the status quo of leaving such negotiations to the commercial entities as a first pass is ideal, but that the FCC could become the final arbiter when there are breakdowns in such negotiations. Such settlement based interconnections between ILECs, CLECs, and IXCs are not that different from interconnections between ISPs, and it is not difficult to imagine the FCC extending their recommendation from access charges and reciprocal compensation to ISP interconnections.

If government regulation of interconnections is deemed necessary, as seems likely based on the discussion above, the issue of settlement-based vs. settlement free interconnections will need to be addressed. At this point, there are too many issues with settlement based interconnections such as defining a valid accounting unit that can be measured with today's equipment and the fact that it is very easy to generate bogus traffic that is nearly impossible to detect and tract that could shift the revenue flows. Therefore, it is unlikely that settlements can be a part of any government regulation for the foreseeable future. Without such a mechanism, deciding which networks should connect to each other on a peer basis (i.e. no money changes hands) and on a transit relationship (i.e. one network pays the other) is difficult to do. It simply cannot be that any network can freely interconnect to any other network, or a small regional network would get a "free ride" on the infrastructure of a large national network. Some kind of classification system, similar to the one defined by the Brokered Private Peering Group will need to be administrated by the government.

It is obvious that the regulation of Internet interconnections is not a trivial manner and will take a very large effort by the FCC and the Internet industry. It is also obvious that any such work will not please all parties involved. However, as the Internet becomes the transport medium for all kinds of services, including such vital elements as 911 emergency connectivity, ensuring end-to-end reachability for all connected entities will certainly become necessary. The only question that remains is how quickly the Internet will become vital for enough services that relying on market pressure to ensure reachability is no longer enough. It may take one prolonged "outage" caused by interconnection negotiations between two networks gone bad, such that a large customer or customers lose enough money or, potentially something even more tragic such as loss of life or property as more and more crucial services move to the net, before the government steps in.

Appendix: Timeline

1890 Congress passes Sherman Antitrust Act

1914 Sherman Act amended by Clayton Antitrust Act

1934 Communications Act of 1934 passed

1964 Paul Baran of RAND publishes "On Distributed Communications Networks"

1966 FCC opens Computer Inquiry I

1969 Experimental ARPANet goes on-line with funding from DARPA

1971 FCC issues final decision for Computer Inquiry I declaring that data services would not be regulated and that common carriers could only offer data services through affiliates

1973 Vint Cerf develops TCP/IP

1976 ARPANet begins experimenting with TCP/IP

1980 FCC issues final ruling for Computer Inquiry II, defining basic service provider and enhanced service provider and that ESPs would not be regulated

1983 All computers on ARPANet required to use TCP/IP

1984 MILNet separate from ARPANet

1984 NSF creates five super computer centers via NSF backbone, connects to ARPANET

1985 NSFNet created with multiple regional networks and university computer centers

1986 FCC released Computer Inquiry III, replacing the "separate affiliate" stipulation of Computer Inquiry I with a model of safeguards

1987 NSF awards contract to Merit, in partnership with IBM and MCI, to upgrade the NSFNet to T1 speeds

1989 NSFNet upgrade to multiple redundant T1's between nodes

1989 Merit, IBM, and MCI present plan to upgrade network to T3

1990 Merit, IBM, and MCI create Advanced Network Services (ANS) to operate the backbone

1990 NSF and Congress sponsor "Commercialization and Privatization" workshop

1991 Upgrade of NSFNet to T3 completed

1991 Commercial Internet eXchange (CIX) and Federal Internet eXchange (FIX) established

1991 Metropolitan Area Exchanges (MAE's) established

1992 Congress subcommittee on Science, Research, and Technology conducts investigation of Privatization

1993 NSF issues solicitation NSF 93-52 to breakup the NSF backbone into a multiple backbone architecture with interconnections at Network Access Points (NAPs)

1995 NSFNet officially shut down

1996 Telecommunications Act of 1996 passed

1997 Emergence of hybrid private-peering and private peering between large networks

1997 UUNET and Sprint announce end of "free" peering unless the peer met certain requirements

1997 FCC issues Access Reform Order declaring ISPs were not required to pay interstate access charges for local Internet connection calls

1998 FCC declares that Information Service Providers would not have to contribute to the Universal Service Fund

1998 Mike Gaddis introduces industry self-regulation proposal of Brokered Private Peering

1999 FCC releases Reciprocal Compensation Order declaring that common carriers were bound by their existing reciprocal compensation obligations to ISPs

1999 Internet Freedom Act of 1999 introduced by Senator Boucher and Goodlatte to lift Inter-LATA restrictions on RBOCs

1999 FCC issues order revising the Communications Act of 1996's original order of unbundling network elements, declaring that ILECs were not required to unbundle certain network elements for broadband Internet connections

References

Baran, Paul. "Publications in the On Distributed Communications Series." URL: http://www.rand.org/publications/RM/baran.list.html.

Baran, Paul. "On Distributed Communications Networks." RAND memorandum RM-3420-PR. August, 1964. Available at URL: http://www.rand.org/publications/RM/RM3420/.

Barrett, Randy. "UUNET Sets Official Peering Policy." Inter @ctive Week, May 13, 1997. Available at URL http://www.zdnet.com:80/intweek/daily/970513a.html.

Blumenfield and Cohen Law and Technology Group. "Telecom Industry History." Available at URL: http://www.technologylaw.com/telephony.html.

Brock, Gerald W. 1994. "Telecommunications Policy for the Information Age." Cambridge, Mass. Harvard University Press.

Cannon, Robert. "Legal Developments: FCC releases Reciprocal Compensation Order. Boardwatch Magazine. May 1999.

Cannon, Robert. "Legal Developments: FCC Finds Broadband Deployment to be 'Reasonable and Timely.'" Boardwatch Magazine. April 1999.

Cannon, Robert. "Legal Developments: Congress Addresses Internet Access, Spam, and Cyberporn." Boardwatch Magazine, July 1999.

Cook, Gordon. "Peering and Transit Issues NANOG Discussion Illuminates Some Current Practices and ISP Concerns." The Cook Report on the Internet. July, 1996.

Cook, Gordon. "UUNET and Sprint Ending Free Peering." The COOK Report on the Internet, Extra Edition. May 3, 1997.

Cook, Gordon. "David Holub Explains Model of an Open Peering, Open Interconnect, 'Bill and Keep' Internet." The Cook Report on the Internet. July, 1997.

Cook, Gordon. IBM Global Net takes Advantage of v-net reach to become $19.95 dial up global ISP. Cook Report, July 1998.

Cukier, Kenneth. "Peering and Fearing: ISP Interconnection and Regulatory Issues." URL: http://ksgwww.harvard.edu/iip/iicompol/Papers/Cukier.html

Depalma, Jennifer. Maturation in a Fee Market: The Changing Dynamics of Peering in the ISP Industry. Forthcoming from the Cato Institute under Telecommunications Policy Analysis: http://www.cato.org/research/telecom-st.html.

Dunlap, Charlotte. "UUNET Defends Peer Pricing." CMPs TechWeb. May 2, 1997. Available at URL http://www.techweb.com:80/wire/news/may/0502uunet.html.

H**alabi, Bassam**. 1997. Internet Routing Architectures. Cisco Press.

Ellis, John. "The forecast for phone rates: flat." St. Petersburg Times, February 27, 1999.

Fleishman, Glenn. "Border Wars on the Net." June 1997. Available at URL http://www.tidbits.com/tb-issues/TidBITS-383.html#lnk4.

Frieden, Rob. Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization. Virginia Journal of Law and Technology. Fall 1998.

Gaddis, Michael. "Brokered Private Peering." Boardwatch Magazine, May 1998.

Geist, Rudolph. "Policy Forum: 1999 – The Year of the ISP – CLEC. How and Why you Should become a CLEC." Boardwatch Magazine. May 1999.

Geist J., Rudolf. "Policy Forum: Regulation of ISPs Likely in the Foreseeable Future." Boardwatch Magazine. July 1999.

Gibson, Robert C. "The Need for Facilities-Based Internet Backbone Competition." May 6, 1998.

Harris, Susan R., and Gerich, Elise. "Retiring the NSFNET Backbone Service: Chronicling the End of an Era." URL: http://www.merit.edu/nsfnet/retire.connecions.html.

Huston, Geoff. "Interconnection, Peering and Settlements – Part I." Internet Protocol Journal, Volume2, Number 1, March 1999. Soft copy available at URL http://www.cisco.com/ipj

Huston, Geoff. "Interconnection, Peering and Settlements – Part II." Internet Protocol Journal, Volume2, Number 2, June 1999. Soft copy available at URL http://www.cisco.com/ipj

Kahin, B., Editor. Commercialization of the Internet Summary Report, RFC 1192. November 1990. URL: available at http://www.cis.ohio-state.edu/htbin/rfc/rfc1192.html.

McCarthy, Bill. "Editor's Notes." Boardwatch Magazine. June 1999.

Oram, Andy. "ISPs and Internet Policy: The New Agenda." The American Reporter, Vol. 5, No 1037. March 30, 1999. URL available at http://www.american-reporter.com/1037/1.html.

Oxman, Jason. "The FCC and the Unregulation of the Internet." Office of Plans and Policy (OPP), Federal Communications Commission. July, 1999

Norton, William. "Interconnection Strategies for ISPs: The Application of a Neutral Internet Business Exchange."

Perlman, Radia. Interconnections. Addison-Wesley Publishing Company. Massachusetts. 1992.

Rickard, Jack. "Internet Architecture: The Internet – What is it?" Internet Service Providers Directory. Available at URL http://boardwatch.internet.com/isp/internetarch.html

Rickard, Jack. "Yet Another Unique Moment in the Time Peering Redux – Back to the Future and the Essentials of a Competitive Internet." Boardwatch Magazine, Editor's Notes, May 1998.

Tappendorf, Sean. The ARPANET and Beyond. URL Available at: http://clavin.music.uiuc.edu/sean/internet_history.html. July 30, 1995.

Thompson, Jim. "No Bed of Roses for CLECs: Problems Abound but Rewards can be Worth it. A Survey of the CLEC Industry in 1999." Boardwatch Magazine. May 1999.

Werbach, Kevin. "Digital Tornado: The Internet and Telecommunications Policy." Office of Plans and Policy (OPP), Federal Communications Commision. March, 1997.

Zakon, Robert H'obbes. Hobbes' Internet Timeline v4.0. Available at URL: http://info.isoc.org/guest/zakon/Internet/History/HIT.html

ARPANET--U.S. DoD-SPONSORED NETWORK, URL Available at: http://www.chaos.com/learn/History.html

"Notes from a symposium of executives/regulators" Symposium was called "Changing Times, Changing Technology," and was held on April 17, 1999 and was brought together by the Harvard Journal of Law and Technology. Posted to cyber-telecom mailing list.

"About the Commercial Internet Exchange," URL available at http://www.cix.org/CIXInfo/about-cix.html.

FCC. CC Docket 96-262 FCC Access Charge Reform Homepage. 7/16/99. URL http://www.fcc.gov/isp.html.

FCC. Telecommunications Act of 1996. URL: http://www.fcc.gov/telecom.html.

FCC. "Reciprocal Compensation Fact Sheet" URL: http://www.fcc.gov/Bureaus/Common_Carrier/Factsheets/nominute.html.

FCC. Communications Act of 1934. URL http://www.fcc.gov/Reports/1934new.pdf.

FCC. "FCC Promotes Local Telecommunications Competition: Adopts Rules on Unbundling of Network Elements." September 15, 1999. URL: http://www.fcc.gov/Bureaus/Common_Carrier/News_Releases/1999/nrcc9066.html

FCC. "In the Matter of Implementation of the Local Competition Provisions in the Telecommunications Act of 1996. CC Docket No. 96-98" URL http://www.fcc.gov/Bureaus/Common_Carrier/Orders/1996/fcc96483.txt

"Customer View of Internet Service Performance: Measurement, Methodolgoy, and Metrics." Cross-Industry Working Team Internet Protocol Performance Measurements workgroup. URL: http://www.xiwt.org/documents/IPERF-paper.pdf.

MSN Encarta Online. Available at URL http://encarta.msn.com.

NetFlow Services and Applications. Cisco Connection Online (CCO). August, 1999. URL: http://www.cisco.com/warp/public/cc/cisco/mkt/ios/netflow/tech/napps_wp.htm.

Section by Section summary of the Internet Freedom Act of 1999. URL:

http://www.house.gov/goodlatte/sectsum.htm.

"Interview: FCC Chairman William Kennard." Internet Industry Directory. URL: http://www.internetindustry.com/Interviews/kennard.shtml.

NSF. NSF 93-52 – Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNET and the NREN (SM) Program. May 6, 1993. National Science Foundation.

NSF. Program Solicitation NSF 92-52, Questions and Answers. June 23, 1993. National Science Foundation.