PRIVATE NETWORKS PUBLIC OBJECTIVES
E. Noam and A. NíShúilleabháin (Editors)

Towards a Public Metanetwork: Interconnection, Leveraging, and Privatization of Government-Funded Networks in the United States

Brian Kahin and Bruce McConnell

1. THE PARADIGM: PARTIAL PUBLIC FUNDING OF AUTONOMOUS INTERNETWORKS

Over the past twenty years the federal government has played a major role in the development of computer networks, initially in the development of the ARPANET under the Advanced Research Projects Agency of the Department of Defense (now DARPA), and subsequently in the development of the "National Research and Education Network" (NREN). Public funds, federal and state, have partially subsidized the creation of dozens of private, predominantly nonprofit, research networks over the past six years. These networks include state networks, multistate regional networks, metropolitan area networks, networks connecting supercomputer centers and their users, and other special purpose networks. These networks have been supported by a combination of revenue sources -- federal funding, state funding, user fees, and contributions from private industry.

As "enhanced services," computer networks have gone unregulated. Even after the Ninth Circuit determined that the FCC could not preempt state regulation of enhanced services,[1] few states have attempted to do so. However, the funding policies of the federal government and (to a lesser extent) state governments have played a major role in shaping a new computer networking environment. This environment has its roots in the packet-switching technology pioneered by the ARPANET in the 1970s, the nonproprietary TCP/IP protocol suite, the interconnection between the ARPANET and CSNET[2] in 1983 that "created" the Internet, and the NSFNET program launched by NSF in 1985. It has emerged as a decentralized quasi-public infrastructure of autonomous internetworks, no longer limited to the academic research community.

The public funding that supports this infrastructure has been characterized by proliferating policy objectives, including a growing embrace of traditional infrastructure rationales. At the same time, a distaste for government ownership and control has been combined with strategic leveraging of private investment. Despite the presence of public funding, this decentralized environment is not well understood by policy-makers and or even by many of the potential stakeholders. While the telecommunications carriers have supported the NREN legislation,[3] they have been slow to develop Internet-related services.[4] The Internet has been better appreciated and exploited by the computer industry, which has produced much of the driving technology, benefitted from the externalities of networking, and is accustomed to a decentralized multidimensional market environment. Public funds have supported research networking in the interests of increasing communication and cooperation among academia, industry, and government researchers. At the federal level, this funding has been aimed at basic research; but at the state level (where there is little public funding of basic research), the expressed goals often include economic development. State and regional networks have embodied a mix of public, private, and nonprofit investment, and have put forth a broader agenda than NSF and the White House Office of Science and Technology Policy have at the federal level. However, the High Performance Computing Act of 1991 reflected the growing political constituency for the NREN and growing expectations that it would help shape a universal information infrastructure.

The degree of direct public subsidy for research networks varies from zero (BARRNet in the San Francisco Bay Area, NEARnet in New England)[5] to 100% in the case of mission agency internetworks (ESNET, NASA Science Internet). Several networks received NSF subsidy at the outset but have since become self-sufficient. Some are organized as cooperatives; some are projects of private universities;[6] some are private, for-profit companies; some are owned by federal agencies but operated privately. Except for the agency networks, virtually all the networks of the Internet are private[7] and considered autonomous by the National Science Foundation and other federal agencies, even though the network may still receive public funds. Specifically, NSF does not dictate the acceptable use policy or other internal policies of grantee networks.

NSF prescribes an acceptable use policy only for the NSFNET backbone, which it makes available for all "mid-level" networks to use free of charge for qualifying traffic. According to the policy the purposed of the NSFNET backbone is "to support open research and education in and among US research and instructional institutions, plus research arms of for-profit firms when engaged in open scholarly communication and research."[8] By providing the backbone as free transcontinental trunk, NSF has greatly stimulated the growth of the Internet for research and education purposes. However, this policy has had the effect of segregating the commercial and noncommercial sides of the Internet and creating considerable confusion about the nature of the Internet.

In one sense, the noncommercial part is more "public" because traffic over the NSFNET backbone is restricted to the public purpose that is defined in the acceptable use policy (AUP). In another sense, the commercial portion is more public because use of it is virtually unrestricted, like the public voice network and the conventional public data networks (SprintNet and Tymnet). As we shall see, this division along with the NSF's cooperative agreement with Merit for the operation of the NSFNET has led to considerable controversy over the commercialization of the Internet and, in 1992, to hearings before the Subcommittee on Science of the House Committee on Science, Space and Technology.

2. CHARACTERISTICS OF THE INTERNET

The conceptual difficulties surrounding public and private characteristics are compounded by the technological characteristics of the network, which differ radically not only from the voice network but from more familiar fully private, centralized, or single-function computer networks. Indeed, the Internet is bedeviled by stereotypes. It is often perceived simply as an electronic mail network which public funding makes available free to universities.

In fact, the Internet is a generalized, multi-function, multi-protocol infrastructure, leveraged from three directions: 1.) enormous private investment in computers and internal networks; 2.) carrier investment in local and long distance fiber optic networks driven by the market for conventional voice and fax and special purpose data networks, as well as anticipated future demand for video services; and 3.) federal and private investment in very high-speed networking technologies, including federal support for high-end applications such as remote visualization.

These different leveraging forces give the Internet its unique character. 1.) It is a bottom-up *collaborative* enterprise, driven by the aggregation of demand at the institutional or campus level. Much of the technical design and development has been the product of volunteers working through the ad hoc Internet Engineering Task Force.[9] 2.) It is comprised of *overlay networks*, riding on leased lines and requiring minimal capital investment; its computer-based routing technology has benefited from rapid advances in microprocessor technology and the growth of LANs. 3.) Funding for high-end uses, statistical multiplexing, and low marginal costs for additional capacity result in *capacity-based pricing* at the wholesale level. Costs are allocated to end users at low flat rates or are simply absorbed into overhead at one level or another. Hence, the common but erroneous perception that use of the Internet is free.

The Internet is not a network but a *metanetwork* -- a functionally defined international set of interconnected, inter-operating, but autonomous networks. While the global telephone network is a homogeneous user-to-user metanetwork of 4 kilohertz voice channels, the Internet is a heterogeneous multi-object metanetwork of indeterminate links, which connect users to computers and information resources as well as to other users. The Internet extends to files that are readily addressable from the network, as well as users addressable from the network. It is defined by the interaction among several different types of "network objects" -- computers, applications, databases, other networks, as well as users and lists of users (enterprises).[10] In principle, each site on the Internet is accessible to any site on any interconnected network for a standard set of functions that includes file transfer, remote log-in, and electronic mail. Originally, the Internet was defined by the TCP/IP protocol suite, but protocol conversion makes it possible to include sites accessed through OSI and DECNET protocols. A user may gateway from one computer into another computer or another network, either transparently or by explicit direction, and the final destination may not be on the Internet in the usual sense of being directly addressable from the Internet.

Other definitional problems arise from the one-way nature of certain links. Corporate networks may be set up so that company users can Telnet and FTP out from the corporate network, but other Internet sites cannot FTP and Telnet into the corporate network. There is a more inclusive metanetwork, sometimes called the "Matrix,"[11] of more limited functionality. Sites in the Matrix, although not necessarily capable of real-time inter-operation through FTP, Telnet, and similar protocols, are nonetheless able to exchange electronic mail through common addressing and routing conventions. The Matrix includes a number of commercial services with mail gateways into the Internet -- MCIMAIL, AT&T Mail, Prodigy, and CompuServe -- as well as the cooperative dial-up networks, such as FidoNet. But because real-time interoperation is not required for mail, compatibility requirements are minimal. There is no central registry of addresses or institutional focus for the Matrix.[12]

There are a wide variety of unsubsidized commercial and noncommercial electronic mail services. There are also commercially available X.25 public data networks, such as SprintNet and Tymnet. However, the TCP/IP networks are distinct in that they support multiple services. They offer fast file transfer using File Transfer Protocol (FTP), which is the largest traffic component on the NSFNET backbone. (See table.) Using the TELNET protocol, they typically support much higher speeds of remote login and interactive use than the X.25 PDNs. (Most universities are connected at T-1 speeds (1.544 mbps); even small institutions are connected at 56 kpbs.) FTP, TELNET, and mail (SMTP) services are all available in a common network environment.

NSFNET T1 Backbone Traffic Distribution by Service
March 1992
Packet Total: 8,924,601,550
Byte Total: 1,641,789,183,650

Service Name	Port	Packet Count	% Pkts	Byte Count	% Bytes
ftp-data	20	2,345,582,650	26.282	828,850,604,950	50.485
(other ports)	--	1,819,683,100	20.390	171,983,064,750	10.475
telnet	23	1,109,624,000	12.433	78,901,539,450	4.806
nntp	119	930,703,700	10.429	178,094,436,150	10.848
smtp	25	779,172,450	8.731	118,713,316,000	7.231
nameserver	53	616,370,700	6.906	62,653,762,550	3.816
ftp	21	210,861,050	2.363	15,199,203,900	0.926
(other protocols)	-	187,167,800	2.097	15,913,880,450	0.969
vmnet	175	179,897,400	2.016	65,039,534,550	3.962
irc	6667	124,715,150	1.397	10,249,531,050	0.624
who/login	513	80,673,900	0.904	3,478,657,300	0.212
ntp	123	66,563,250	0.746	5,078,830,450	0.309
(unknown)	1023	64,405,950	0.722	9,733,875,200	0.593
X0	6000	47,513,250	0.532	7,010,593,600	0.427
syslog/cmd	514	39,135,050	0.439	16,412,937,450	1.000
talk	517	35,336,300	0.396	3,623,556,450	0.221
finger	79	34,945,600	0.392	3,528,747,400	0.215
snmp	161	22,898,450	0.257	2,605,795,550	0.159
(unknown)	1022	22,293,750	0.250	3,844,914,900	0.234
uucp	540	17,925,000	0.201	3,929,692,550	0.239
NSS Routing	159	17,281,650	0.194	1,413,335,700	0.086
src	200	13,387,100	0.150	783,693,150	0.048
(unknown)	1021	12,595,900	0.141	2,496,509,800	0.152
ntalk	518	8,396,300	0.094	902,138,400	0.055
(unknown)	1020	7,640,950	0.086	2,063,678,750	0.126
tftp	69	6,778,750	0.076	378,432,000	0.023
router efs	520	6,282,700	0.070	3,003,212,200	0.183
mdec net	700	6,127,900	0.069	1,509,414,900	0.092

Source: Merit, file= <NIC.MERIT.EDU> /nsfnet/statistics/1992/t1-9203.ports

NSFNET Backbone Traffic Distribution by Service
June 1994
Packet Total: 75,280,338,950
Byte Total: 15,350,480,259,500

Service Name	Port	Packet Count	% Pkts	Byte Count	% Byts
ftp-data	20	14,920,693,000	19.820	5,406,765,553,400	35.222
(other_tcp/udp_ports)	999	12,136,852,350	16.122	1,802,547,568,950	11.743
telnet	23	10,102,897,900	13.420	742,523,290,550	4.837
nntp	119	7,573,275,950	10.060	1,668,585,813,300	10.870
smtp	25	5,851,752,650	7.773	988,020,772,400	6.436
ip	-4	4,014,654,250	5.333	1,204,497,949,900	7.847
domain	53	3,865,035,100	5.134	374,096,786,350	2.437
www	80	3,060,869,850	4.066	946,538,669,550	6.166
icmp	-1	2,987,323,650	3.968	295,580,874,850	1.926
gopher	70	2,012,045,400	2.673	567,478,655,950	3.697
irc	6667	1,897,549,600	2.251	193,334,424,000	1.259
ftp	21	1,389,715,300	1.846	129,841,696,950	0.846
talk	517	889,821,500	1.182	92,412,534,800	0.602
X0	6000	758,177,850	1.007	162,151,234,500	1.056
login/who	513	381,036,350	0.506	43,481,370,550	0.283
vmnet	175	351,185,300	0.467	142,147,311,950	0.926
(unknown)	1023	331,001,650	0.440	35,440,965,350	0.231
snmp	161	214,356,200	0.285	24,753,750,800	0.161
unidata-ldm	388	197,048,450	0.262	60,324,898,600	0.393
ntp	123	174,574,100	0.232	13,000,123,600	0.085
finger	79	163,632,900	0.217	20,579,482,950	0.134

Source: ftp://nis.nsf.net/statistics/nsfnet/1994/

The TCP/IP platform supports a wide variety of other services. See Table 1. showing the most widely used services on the NSFNET T1 backbone. There are nearly 200 "Well Known" services identified to ports 1 to 256. An additional like number are documented but used inconsistently across the Internet.[13] And there are unknown implementations using unassigned TCP ports (the sum of the "unknowns" equals the figure for "other ports"). Many of these services are experimental, but it is also possible to set up a virtual private network using common services on unassigned ports.

In effect, the backbone provides an open platform for the implementation of a wide variety of public and private networks supporting any combination of functional services in either an open or private manner. In this sense, it offers both a service platform for a wide number of applications in support of research, such as remote visualization, and a technology platform which offers opportunities for experimentation in computer-to-computer communication. In implementing a T-3 (45 mbps) backbone 1991-92, NSF pushed a particular technological envelope directly. On the other hand, the mail traversing the NSFNET backbone has probably had the largest impact on the greatest number of people. The multifaceted impact of the federal investment in research networking came to the fore in the rhetoric and debate that attended the high performance computing legislation introduced by Senator Gore.[14] The language on the NREN in the enacted High Performance Computing Act of 1991 reflects the variety of constituencies that took interest in the bill -- and the diversity of policy objectives that were read into it. These diverse but related objectives combine with the protean nature of the rapidly growing Internet to create a new policy environment where interaction among the sectors is accelerating dizzingly -- and where the tilt, size, and shape of the playing field are relative to perspective and the time of day. Before parsing out the policy objectives in plans for an NREN, however, we briefly review the history of federal funding for networks, beginning with its own voice network.

3. FTS TO FTS 2000 -- FROM NETWORKS TO NETWORK SERVICES

After the Second World War, Federal agencies individually satisfied their long distance telecommunications requirements by leasing circuits and services on an ad hoc basis from the Bell System. In 1961, traffic overloads during the Cuban Missile Crisis revealed weaknesses in the Federal government's communications systems and stimulated the development of a private, government-operated long distance telephone network, the Federal Telecommunications System (FTS). By 1980, AT&T supplied the Federal government, through contracts with the General Services Administration (GSA), a dedicated voice-grade network of 15,000 network trunks and fifty-two switching nodes. The network served 1.3 million Federal users at over 4,000 locations.[15]

Deregulation and rapid changes in telecommunications technology eliminated the cost advantages of FTS and strained GSA's limited technical resources. During the mid-1980s, GSA decided to replace the FTS with new contracts. Instead of a dedicated voice network, however, FTS 2000 would be a voice and data service offering, with pricing on a service usage basis. The decision to pursue a service offering instead of engineering a Federal network was grounded in practical and policy realities. GSA realized that it would be unable to keep pace with the rapidly changing technical and market environment. A government-built network would always lag behind the best private sector offerings. Strong pressure within the Reagan Administration to privatize government operational activities favored contracting out of network engineering and operations. Finally, key Congressional figures advocated Federal support of emerging competition in the provision of public switched network services, and for increased competition in contracting generally.[16]

GSA planned a winner-take-all competition in which the winner would be entirely responsible for engineering and operations, leaving GSA to manage the contract on behalf of all Federal agency users. By the time GSA had a request for proposals on the street in early 1987, however, only three firms (AT&T, MCI, and Sprint) remained which provided nation-wide telecommunications service. Although all three announced they would compete for the award, uncertainty in both the technical and business requirements of the government led Sprint to drop out of the competition in mid-1987. To increase the amount of competition (and thus increase pressure on the vendors to lower prices), the solicitation was restructured to provide for a two-vendor award, Sprint re-entered the competition, and along with AT&T, was awarded a ten-year contract in 1988 to provide voice and data telecommunications services to Federal agencies. By law, FTS 2000 is mandatory for Federal agency use, except for non-administrative Defense Department traffic.[17] Telecommunications between Federal contractors and grantees and their sponsoring agencies are generally not provided under the FTS 2000 contracts because these services are generally acquired under grant or cooperative agreements not subject to GSA's authority.[18] FTS 2000 is a private network in that it is privately owned, provided under specific contractual arrangements, and is accessible only to government officials or contractors conducting government business. Local access is provided by local phone companies and the regional Bells under subcontracts with AT&T and Sprint. Sprint's inter-LATA services are provided for entirely over its part of the Public Switched Network, whereas AT&T has installed dedicated switches to serve its FTS 2000 customers.

4. ARPANET -- A VERY DIFFERENT MODEL

Meanwhile, in 1969, the Defense Advanced Research Projects Agency (DARPA) established ARPANET, an experimental network connecting Defense agencies with their contractors and grantees. The network was designed to demonstrate the potential of computer networking made possible by packet-switching technology and the resource sharing it allowed. The DARPA-developed packet-switching technology was later successfully commercialized by private sector firms such as the Tymnet and Telenet services. DARPA sponsored several additional networks during the 1970s and developed the Internet protocols, which allowed messages to be interchanged across multiple interconnected networks. By 1983 ARPANET became congested and Defense split its military R&D traffic onto a separate network, MILNET. Defense shifted responsibility for civilian R&D networking to the National Science Foundation in 1985, and shut down the ARPANET completely in 1988.[19] The ARPANET experiment differed from FTS 2000 in several ways. First, it was an experimental network, not a production network. While FTS 2000 tried to give government users access to the most advanced commercial offerings, DARPA used its network to advance technology beyond existing offerings. Its experimental nature was recognized by the research community, which was thus willing to tolerate its idiosyncratic performance and arcane user interface. In accessibility and usability, ARPANET was more restricted than FTS/FTS 2000, although, unlike FTS it could be used by those outside the government. Second, ARPANET was a collaborative Federal-university-industry development effort, not a arms-length competition among private firms wishing to supply commercial services to a large customer.

Consequently, a large portion of the development funding came from the R&D budgets of the computer and communications firms involved in the collaboration.

5. CSNET

By the late 1970s the contribution of the ARPANET to research collaboration became apparent throughout the US computer science community. However the ARPANET connected only those university sites such as Stanford, Berkeley, and MIT, which had DARPA contracts. Computer science researchers elsewhere felt that the network was producing an unfortunate situation of haves and have-nots in terms of connectivity. In 1979, some university computer science researchers, including Larry Landweber, Dave Farber, and Tony Hearn, collaborated on a proposal to NSF to extend ARPANET technology to other university sites. After several rounds with the NSF Science Board, and with the assistance of Rick Adrian at NSF, a program was agreed to. The new network, dubbed CSNET, was to become self-supporting in five years. ARPANET protocols (both X.25 and what was then NCP/IP) would be used throughout, except for dial-up connections which only offered e-mail.[20] To achieve self-sufficiency, dial-up was charged on a per-minute basis, and dedicated access on a capacity basis.[21] A management committee, and later a university corporation,[22] managed the network as not-for-profit business. Operations were handled by BBN, and a memorandum of understanding was developed between NSF and DARPA to allow the exchange of traffic. The first CSNET sites came up in 1982, with 40 sites operational in 1983. In 1988, 200 sites were connected in 12 countries, including commercial sites where the business traffic was related to Federal research.

An important effect of CSNET was to train many people in internetworking technologies, sowing the seeds for the eventual development of today's regional network structure. Federal funding, which began in 1981 and ended in 1986, went principally for technology development and the acquisition of four VAX computers to serve the network. As the NSFNET began to develop, CSNET ceased to grow and eventually terminated in 1991.

6. NSFNET

In 1985, the NSF began to fund the development of the NSFNET, a national high-speed backbone with connections to six supercomputing centers, five of which had also just been established by NSF. Beginning in 1987, NSF funded the backbone network under a cooperative agreement with Merit, Inc., a non-profit arm of the University of Michigan, to operate the backbone. IBM and MCI contributed switching and transport technology in a cost-sharing arrangement. NSF also established two additional funding programs to support the NSFNET. The first was grants to networks, the first of which were organized around supercomputer centers. Others were designed to serve states or regions. Some of these became projects of existing cooperative organizations, like SURANET, formed by the Southeastern Universities Research Association. NYSERNet was created by the NSF-funded Cornell Supercomputing Facility to allow it to share its computing resources with other research institutions in New York state, as such it became a state initiative supported by the legislature.[23] SURANET covered fourteen states, while four networks made their home in California. Some regional networks developed on their own, including a few for-profit networks that connected commercial firms to the Internet. Although NSF sought to ensure that the entire country was served by research networks, none of the networks enjoyed any kind of exclusive monopoly, and some found themselves competing with each other, especially on their outer fringes.

The second additional network funding NSF provided was the "Connections" program, which offered grants of $20,000 to cover the costs of switches (called "routers") and initial dedicated circuit connection from universities to the regional networks. The Table below shows the funding of these programs during fiscal years 1988-91:

NSFNET BUDGET ($ millions)

Fiscal Year	1988	1989	1990	1991
Backbone	2.7	3.6	4.9	8.8
Regional Grants	3.6	7.1	7.1	7.3
Connections	0.2	0.5	0.7	1.4
Other	1.2	2.3	2.4	3.5
Total	7.8	13.6	15.0	21.0

Source: National Science Foundation

The contribution of Federal funds toward the total costs of developing and operating the NSFNET is variously estimated from 1/3 to 1/40, depending on how the NSFNET is defined, on the case that is being made, and what is included in the costs. Part of the reason for this high degree of leveraging is that the results of Federal research efforts in networking technologies can be transferred to the private sector faster and more completely than in some other research areas. Networking research is more like biotechnology research than like high energy physics. Access is much broader than with the ARPANET, for two reasons. First, as a part of the Internet, the NSFNET provides carriage to Internet users based on reciprocal agreements with other Internet networks. Second, the growing number of users on the NSFNET (currently over 1000 institutions are connected) mean that its "privateness" is lost.... A watershed decision during the mid-1980's was NSF's choice of the TCP/IP protocol rather than a proprietary protocol or X.25. As Mandelbaum and Mandelbaum observe:

> "It led almost directly to the establishment of the system of specialized private academic networks we have today, rather than to reliance by the academic and research community on the public, commercial networks that are the mainstays of the business world."[24]

7. MISSION AGENCY RESEARCH NETWORKS

As NSFNET developed, several Federal research agencies (principally NASA and the Energy Department) saw a need for reliable, uninterrupted IP connectivity with their principal research partners. These networks are interconnected with the NSFNET but are controlled directly by the mission agencies and dedicated to mission requirements. The autonomy of this arrangement allows the agencies to shed or delay unessential messages during a critical mission such as a space shot. These networks are more like the original FTS -- dedicated, government-owned networks, although they may be operated by contractors. In addition, they are primarily production networks supporting ongoing research programs, not experiments in network technology.

8. INTERAGENCY INTERIM NREN

The National Research and Education Network was first described in 1989 as providing "high-speed communication access to over 1300 institutions across the United States within five years. . . . so that the physical distance between institutions is no longer a barrier to effective collaboration."[25] The concept had evolved (and continues to do so) from the national research network that the NSFNET represented. In 1989, it became one element of the President's High Performance Computing Program, which was renamed the following year as the High Performance Computing *and Communications* Program, in recognition of the importance of the network component. The HPCC is a Federal R&D effort made up of four elements, and is "designed to sustain and extend US leadership in all advanced areas of computing and networking."[26] A key issue for the HPCC initiative is the definition of the migration path from the existing NSFNET to the NREN. This migration is complicated by a lack of agreement among the principal constituencies of the NREN as to what the NREN actually is. Officially, it is "the future realization of an interconnected gigabit computer network system supporting HPCC--a network for research and education, not general purpose communication."[27] But there are expectations of a broader agenda. For example, the Computer Systems Policy Project, an affiliation of major US computer systems companies, sees the NREN as "the foundation for something broader and more exciting. . . . that will provide all Americans with access to unique resources, public and private databases, and other individuals throughout the country." CSPP recommends "expanding the activities under the NREN" initiative, and:

> "expanding the current vision of the HPCC initiative to include Grand Challenges motivated by social and economic needs in areas of interest to the government and general public, such as advances in the delivery of health care and services for senior citizens; improvements in education and opportunities for lifelong learning; enhanced industrial design and intelligent manufacturing technologies; and broad access to public and private databases, electronic mail and other unique resources.[28]

In recognition of the evolving understanding of the NREN's ultimate configuration and scope, the Administration in 1991 introduced the concept of the Interagency Interim NREN (IINREN) as an evolutionary step towards the gigabit NREN. The IINREN is an effort to upgrade the NSFNET backbone, assist regional networks to upgrade facilities, capacity, and bandwidth, and to interconnect the backbone networks of other agencies.[29] But the IINREN and NREN are elusive terms because, as for the NSFNET, there is no clear definition of their scope. Indeed, it is difficult to discern and define any institutional structures other than the individual autonomous networks. There is no central management for the NSFNET, the Internet, the IINREN, or the NREN. Language in the High Performance Computing Act specifying NSF as the lead agency for the deployment of the NREN had to be excised, despite the support of the academic community, because the Administration did not want its

HPCC program locked into a particular managing agency. Passage of the bill was threatened repeatedly by the interests of national laboratories within the Department of Energy seeking a leading role in the NREN and other components of HPCA. In the end, the HPCA more or less restates what the agencies are already and adds general goals without specifying how they are to be achieved.

Even the NSFNET, which has been in existence over six years, remains vaguely defined -- an ambiguous collection of networks. The NSFNET backbone is a special case, since it is provided under an a formal agreement and serves as a free transcontinental transport for the 25 or so networks that are authorized to use it for qualifying purposes. Beyond the backbone, it is not clear what the NSFNET includes. NSF only funds some of the mid-level networks connecting to the backbone, but it does not deny access to the backbone for any internetwork wishing to use in accordance with NSF's purpose and policies. Thus, NEARnet and PREPnet (Pennsylvania) were permitted to use the backbone although they did not receive direct funding from NSF.[30] PSInet and AlterNet were also authorized to the NSFNET backbone for qualifying traffic, although they are for-profit services with their own transcontinental backbones.[31] In effect, the NSFNET backbone has served as a wholesale carrier available without discrimination to any internetwork service -- for traffic that satisfied the purpose defined by NSF. Originally this purpose was to support "scientific research and other scholarly activities," then in 1990, "academic research and education," and, finally, in early 1992, "open research and nonprofit instruction" (see Appendix). The nature of the use is determinative, not the status of the user.

With this access to the backbone for qualifying uses, and their own transcontinental backbones for non-qualifying uses, AlterNet and PSI (Performance Systems International) were able to interest companies that were familiar with the Internet in an academic context but would also use it for proprietary research or other purposes. These for-profit services naturally focused their resources metropolitan areas with concentrations of high-tech industry -- San Francisco, Washington, DC, Boston.... They were joined in California by CERFnet, a project of General Atomics which had initially received NSF funding. In 1991, the three services formed an association, the Commercial Internet Exchange or CIX (pronounced "kicks"), which maintained a physical exchange of the same name in the Bay Area. The CIX, as an exchange for sharing traffic, was modeled after the Federal Internet Exchanges (FIX) (of which there were two, one on each coast) which connect the NSFNET backbone and the mission agency internetworks.

In addition to direct competition among network services, for-profit and nonprofit, there is pressure from resale of network connectivity. At one level, this problem is present in university-developed industrial parks, which may offer connections to the university's campus area network. At another level, guest accounts on university computers have long been given out freely, usually without charge. Mid-level networks initially resisted "resale" of networks, at least for corporate connections. The mid-level network's interest in maintaining its revenue base does not hold very well against efficiency and equity arguments (or common carriage principles), especially since connections are priced on capacity. It may be desirable for a state education network to connect as a single customer to a mid-level network -- or even to do so indirectly through a university.[32] Competition among commercial providers now appears to be encouraging a growing number of re-sellers. One example is the World, dial-up host in Boston which retails full connectivity to the commercial Internet through AlterNet. Internet access through the World costs $20/month for 20 hours plus $1/hr. for additional hours. Another example is Netcom, which began as a single host in the Bay Area but now has its own network and can be reached with a local telephone call in most California metropolitan areas and a number of major cities nationwide. Netcom offers *unlimited* dial-up access for $17.50/month. Note that these providers cannot offer interactive services (FTP, Telnet, etc.) across the NSFNET backbone because of the NSF's acceptable use policy (as opposed to mail, which is transported freely throughout the Internet without regard to origination or destination).[33]

The pricing of the unsubsidized commercial Internet services is worth noting because it suggests extremely low costs of operation. However, it must be emphasized that these are dial-up services (up to 9600 baud) and that occasionally all ports may be busy. Furthermore, at the lowest level of service, the user is only on the Internet as an account at the provider's address. A higher level of service is available using SLIP, Serial Line Internet Protocol. Instead of accessing the Internet by logging into a host machine, SLIP service enables the user's own microcomputer to act as an Internet host with its own Internet address. When connected through SLIP, the user can pull remote files directly, instead of transferring them first to the host and then downloading them into the user's microcomputer. For an additional fee, a user may secure a dedicated port on the host system, which assures on-demand access to the Internet. By leasing a line and dedicating equipment on a 24-hour basis, the user can provide on-demand access in the other direction -- i.e., from the Internet -- which is necessary if employees, agents, customers, or suppliers need to access the system at will.[34]

As demand grows and TCP/IP technology is commercialized, mid-level networks themselves are increasingly providing low cost services directly to individual users and small businesses. Thus there may be competition between networks and their own commercial customers, at least in the more competitive markets. For the nonprofit networks, the big stumbling block remains the acceptable use policy which while not binding on them has left them the product of a specific culture, clientele, and charter in a complex, unregulated, increasingly environment.

9. ANS AND THE "PRIVATIZATION" OF NSFNET

The hierarchical model, which has remained essential intact in the telephone world and in which the regional networks once had a clearly defined place, is collapsing -- even at the highest level. Since September of 1990, the NSFNET backbone has no longer existed as a distinct network but as a contract for services from Advanced Network and Services, Inc. (ANS). ANS operates the backbone as part of its own private network, ANSnet, which competes with other network services, commercial and nonprofit, in offering T-1 connections to institutional and corporate networks. ANS was formed as nonprofit corporation by IBM, MCI, and Merit and capitalized by a commitment of $5 million each from IBM and MCI over a three-year period plus $1 million from the State of Michigan Strategic Fund. It became a subcontractor for Merit who continued to hold the 5-year cooperative agreement with NSF that it had won in 1987. It provides network services to Merit, and in turn subcontracts back to Merit for the operation of the NSFNET Network Operations Center in Ann Arbor. This arrangement replaced the joint study agreements that Merit had with IBM and MCI directly in support of the original cooperative. IBM and MCI still contribute technology (as well as funds for ANS), but their contribution has been through ANS. The restructuring of the cooperative agreement with ANS proceeded with NSF's approval as plans were being made to upgrade the backbone from T1 to T3 speeds nearly three years ahead of schedule. This pushing of the technological envelope in a fully operational production network with "cost-sharing" from the private participants was consonant with the evolving federal strategy for a National Research and Education Network. NSF's annual bill for the backbone went from $3 million to $8 million, but its cost for capacity in mbps dropped twelve-fold at the margin.

The advent of ANS was heralded as a unique partnership in the national interest and the new arrangement was described by its President, Al Weis, a former Vice President of IBM, as the privatization of the backbone. The expectation was that restricted NSFNET traffic could be aggregated with additional traffic from companies seeking to market services to users in higher education.[35] It was hoped that commercial users would pay at least the average costs of using the network, while noncommercial use would remain subsidized -- or that noncommercial use could be priced at the margin, while commercial users covered the principal fixed costs.

The additional traffic to be carried by ANSnet was initially restricted by ANS's charter which was intended to qualify it as a 501(c)(3) charitable organization. This left ANS with limited opportunities, so it formed a for-profit subsidiary, CO+RE, Inc. (meaning "commercial" plus "research and education" but pronounced "core"), in May, 1991, as a vehicle for marketing unrestricted commercial use of ANSnet. This put ANS in direct competition (at least for T-1 connections) with the commercial Internet providers, PSI, AlterNet, and CERFnet, who had just formed the Commercial Internet Exchange. Along with the CO+RE subsidiary, ANS announced a plan for commercial services designed to involved the nonprofit mid-level networks connected to the backbone. The plan invited the mid-level networks to hook up commercial customers, either directly or through ANS CO+RE, with the expectation that these commercial customers would be contribute to "the network infrastructure," specifically ANSnet and, through a special fund, the mid-level networks. The plan included measuring traffic to and from declared commercial sites in "COMBits," an original unit designed as a compromise between bits and packets. Although the plan was developed in conjunction with many of the mid-level networks, its complexity was daunting. It presented the networks with three options -- a "Connectivity Agreement," a "Gateway Attachment Agreement," and a "Cooperative Agreement"[36] -- which challenged the networks to make difficult decisions about how commercial their operations would become. While the measurement of COMBits was intended only to provide a guideline to setting annual fees in advance, it was perceived as introducing metering within an environment where capacity-pricing was the established norm and usage-pricing was anathema for both practical and philosophic reasons.

The commercial IP networks, who had just formed the CIX, saw the ANS plan as an effort to dominate the commercialization of the Internet under a centralized, privately controlled model. These commercial providers were all spinoffs of non-profit network organizations in some manner,[37] and so shared with the mid-levels similar roots and perspectives, despite the difference in profit orientation. Although ANS was nonprofit, its self-perpetuating board made it appear more private than most of the mid-levels, most of whom had "members" and some degree of accountability to major institutional users. As a large organization with extremely large corporations behind it -- and as a direct competitor at the high end that was hiring away many of the highly regarded individuals in networking community, ANS was viewed with suspicion by some of the mid-level networks.

The cooperative agreement between NSF and Merit (which in turn subcontracts with ANS) was to expire late in 1992, but in November 1991, NSF announced a new project development plan for the NSFNET. Acknowledging tension between concerns for stability of the NSFNET backbone and competition, the plan announced that in the interests of stability the cooperative agreement would be extended for an additional eighteen months. However, the plan announced three changes for the next solicitation: 1.) The routing authority function would be awarded separately to ensure that tactical advantages would not accrue to a provider of connectivity services; 2.) The awards for connectivity services would not be limited to a backbone model -- so that an award could be made for an exchange or series of exchanges like the CIX and FIX exchanges. (This change was not explicit but was implied in the choice of the term "connectivity.") 3.) There would be at least two awards to provide connectivity. Like the FTS 2000 solicitation, the proposal to enter into at least two cooperative agreements reflected an awareness of the impact of the federal funding on a competitive environment.

The CIX viewed the unwillingness of ANS to join the CIX agreement as evidence of a monopoly position within the NSFNET that extended to providing commercial services through the NSFNET-connected mid-level networks. ANS claimed that the CIX agreement, which required interconnection without compensation or settlements in either direction, would not adequately recognize and compensate ANS for its investment in the T-3 backbone. This debate was clouded by conflicting claims about the T-3 technology and the implications of the cost-sharing upon which the cooperative agreement between NSF and Merit was premised.

These issues were aired in a March 12, 1992, hearing before the House Subcommittee on Science, chaired Representative Boucher. William Schrader, President of PSI, claimed that ANS had favored T-3 technology under development at IBM rather than further developed technology of established TCP/IP router vendors such as Proteon and Cisco.[38] Schrader questioned the fairness of major changes in the cooperative agreement (the ANS subcontract and the substantial increase in funding for the T-3 backbone) without any public process. The NSF Acceptable Use Policy for the backbone was also at issue in the hearings. Schrader and Mitch Kapor, founder and former CEO of Lotus Development Corp. and now Chairman of the CIX, claimed that the AUP operated to give ANS a monopoly over commercial traffic between the sites connected by the NSFNET backbone.

Eric Hood, President of FARNET, the association of (predominantly nonprofit) mid level networks, also testified in opposition to the acceptable use policy and invited Congressional action to remove it. Boucher questioned whether there was in fact a legal basis for requiring an AUP for the backbone.[39] In the end, he successfully sponsored legislation that would explicitly permit NSF to fund facilities for unrestricted use if doing so would increase the value of the network for research and education.[40] However, the arrangement with ANS was left unchanged.

10. CONCLUSION

The Internet continues its dramatic growth. The Internet Society estimates that fifteen million Americans now have full access to the Internet. However, the vast majority of users have access only through institutional or corporate networks. Probably less than 3% have access direct through commercial Internet service providers or public access hosts. This will change significantly over the coming year as the consumer online services, America OnLine and CompuServe, begin to provide fully functional connectivity. Most of the new addresses registered each month are now commercial sites. Many of the regional networks have sold out to private entities or spun off commercial services. The Commercial Internet Exchange now has over 40 members. ANS agreed to interconnect on experimental basis in June of 1992 and finally joined the CIX.

The cooperative agreement between NSF and Merit (and ANS) has been extended a second time and is still in effect, although the transition to the new NSFNET is scheduled for 1994-95. The solicitation was later amended substantially to unbundle connectivity services into three components: a very high-speed backbone (vBNS), which would eventually be limited to experimental uses and specialized research; a number of Network Access Points, which provided access to the vBNS but also permitted the exchange of traffic between commercial networks and "research and education" networks; and grants to mid-level networks to support inter-regional connectivity through the NAPs.[41] In effect, NSF moved toward a policy that, like the regulatory policies of the time, by using unbundling to promote competition rather than dual award compromise of FTS 2000.[42]

THE NSFNET BACKBONE SERVICES ACCEPTABLE USE POLICY GENERAL PRINCIPLE:

(1) NSFNET Backbone services are provided to support open research and education in and among US research and instructional institutions, plus research arms of for-profit firms when engaged in open scholarly communication and research. Use for other purposes is not acceptable.

SPECIFICALLY ACCEPTABLE USES:

(2) Communication with foreign researchers and educators in connection with research or instruction, as long as any network that the foreign user employs for such communication provides reciprocal access to US researchers and educators.

(3) Communication and exchange for professional development, to maintain currency, or to debate issues in a field or subfield of knowledge.

(4) Use for disciplinary-society, university-association, government-advisory, or standards activities related to the user's research and instructional activities.

(5) Use in applying for or administering grants or contracts for research or instruction, but not for other fundraising or public relations activities.

(6) Any other administrative communications or activities in direct support of research and instruction.

(7) Announcements of new products or services for use in research or instruction, but not advertising of any kind.

(8) Any traffic originating from a network of another member agency of the Federal Networking Council if the traffic meets the acceptable use policy of that agency.

(9) Communication incidental to otherwise acceptable use, except for illegal or specifically unacceptable use.

UNACCEPTABLE USES:

(10) Use for-profit activities (consulting for pay, sales or administration of campus stores, sale of tickets to sports events, and so on) or use by for-profit institutions unless covered by the General Principle or as a specifically acceptable use.

(11) Extensive use for private or personal business.

This statement applies to use of the NSFNET Backbone only. NSF expects that connecting networks will formulate their own use policies. The NSF Division of Networking and Communications Research and Infrastructure will resolve any questions about this Policy or its interpretation.

ENDNOTES

[1] *People of the State of California v. FCC*, 905 F.2d 1217 (9th Cir. 1990).

[2] CSNET – Computer Science Network, a multi-protocol network funded by NSF to enable resource sharing among computer science departments and research facilities. CSNET and BITNET were merged administratively in 1988 to form the Corporation for Research and Education Networking (CREN). Both ARPANET and CSNET were defunct by 1991.

[3] The High Performance Computing Coalition is the lobbying group organized to support the legislation introduced by Senator Albert Gore, Jr., beginning in 1988 that eventually became the High Performance Computing Act of 1991.

[4] The RBOCs have been less interested than the IECs, although a few operating companies have participated in state networks, such as NYSERNET and PREPnet, where their visibility has political benefits.

[5] BARRNet was funded by NSF early on but has since become self-sufficient. NEARnet was started without federal subsidy, although at one point it received funding to connect to the new T-3 backbone. Both were acquired by BBN Systems in 1994.

[6] NEARnet was actually a project of MIT with a steering committee that also included Harvard, Boston University, and BBN Systems. However, to receive service one had to pay a "membership" fee.

[7] The occasional exceptions are state networks, such as NYSERNET or CONCERT (North Carolina), which are usually specially chartered by the state legislature.

[8] From 1990 to February 1992, when the current AUP was issued, the stated purpose of the NSFNET was "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." Prior to that, the purpose was "scientific research and other scholarly activities."

[9] The Internet Engineering Task Force and the Internet Research Task Force comprise the Internet Activities Board (IAB). In 1992, the IAB was integrated into the Internet Society as the Internet Architecture Board.

[10] Sophisticated users can leverage network objects with far-reaching consequences. Using *Archie servers* it is possible to locate any file posted using public directory conventions at any site on the Internet. Large files -- be they libelous, pornographic, or pirated -- can be retrieved and broadcast instantly and effortlessly to automated mailing lists of thousands of individuals. By combining such lists into a personal address file, the file can be sent to hundreds of thousands, congesting the network and clogging mailboxes.

[11] John S. Quarterman, *The Matrix*, Digital Press, 1991; John S. Quarterman, "How Big is the Matrix?" *Matrix News*, Vol. 2 No. 2, February 1992.

[12] Except, of course, for the amorphous Internet and its domain style of addressing. Some commercial bulletin boards misleadingly offer access to the Internet when, in fact, they only provide an Internet mailing address.

[13] See "Port Numbers" in Joyce Reynolds and Jan Postel, *Assigned Numbers*, RFC 1060, Internet Activities Board, March 1990.

[14] S. 2918, "The National High-Performance Computer Technology Act of 1988, introduced October 18, 1988. S. 1067, "The National High Performance Computer Technology Act of 1989," introduced May 18, 1989, S. 272, "The High-Performance Computing Act of 1991," introduced January 24, 1991, which became Public Law 102-194, December 9, 1991.

[15] Bennington, Bernard J., "Beyond FTS2000: A Program for Change," Appendix A, National Research Council, Washington, 1989, pp. 1-3.

[16] Kettl, Donald F., "Sharing Power: Public Governance and Private Markets," Brookings Institution, Washington, 1993, pp. 67-98.

[17] Treasury, Postal Service, and General Government Appropriations Act of 1994, P.L. 103-123, Section 620.

[18] The Brooks Act, codified at 41 USC 759, gives the Administrator of General Services exclusive authority to procure computer and telecommunications equipment and services. The law exempts mission-critical Defense Department applications from the GSA authority. A recent Federal appeals court decision reaffirmed the exclusion from the Brooks Act for procurements conducted by Federal labs (US West Communications vs. United States, 940 F.2d 622 (Fed. Cir. 1991)).

[19] Gould, Stephen, "The Federal Research Internet and the National Research and Education Network. Prospects for the 1990s," U.S. Congress, Congressional Research Service, 90-362 SPR, July 26, 1990, pp. 5-6. Reprinted in McClure, Appendix M, pp. 572-3.

[20] IP did not support dial-up at that time, so modem protocols were used underneath the ARPANET protocols.

[21] For a brief period, CSNET used commercial X.25 services, but the cost of hook-ups and per-packet charges proved prohibitive for computer collaboration.

[22] University Corporation for Atmospheric Research (UCAR) in Colorado, which also runs the NSF-funded supercomputer center known as NCAR.

[23] *Ibid*, pp. 60-61.

[24] Mandelbaum, Richard, and Mandelbaum, Paulette A., "The Strategic Future of the Mid-Level Networks," in Kahin, Brian, ed., *Building the Information Infrastructure*, McGraw-Hill, 1992, page 62, fn. 6.

[25] Federal Research Internet Coordinating Committee (FRICC), *Program Plan for the National Research and Education Network*, May 23, 1989.

[26] Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," January, 1992. The four elements support development of hardware, software, networking and human resources.

[27] *Ibid*. Section 102 of the High Performance Computing Act of 1991, offers less of a definition but paints a descriptive picture:
NETWORK CHARACTERISTICS.--The Network shall--
(1) be developed and deployed with the computer, telecommunications, and information industries;
(2) be designed, developed, and operated in collaboration with potential users in government, industry, and research institutions and educational institutions;
(3) be designed, developed, and operated in a manner which fosters and maintains competition and private sector investment in high-speed data networking within the telecommunications industry;

(4) be designed, developed, and operated in a manner which promotes research and development leading to development of commercial data communications and telecommunications standards, whose development will encourage the establishment of privately operated high-speed commercial networks;
(5) be designed and operated so as to ensure the continued application of laws that provide network and information resources security measures, including those that protect copyright and other intellectual property rights, and those that control access to data bases and protect national security;
(6) have accounting mechanisms which allow users or groups of users to be charged for their usage of copyrighted materials available over the Network and, where appropriate and technically feasible, for their usage of the Network;
(7) ensure the interoperability of Federal and non-Federal computer networks, to the extent appropriate, in a way that allows autonomy for each component network;
(8) be developed by purchasing standard commercial transmission and network services from vendors whenever feasible, and by contracting for customized services when not feasible, in order to minimize Federal investment in network hardware;
(9) support research and development of networking software and hardware; and
(10) serve as a test bed for further research and development of high-capacity and high-speed computing networks and demonstrate how advanced computers, high-capacity and high-speed computing networks, and databases can improve the national information infrastructure.

[28] Computer Systems Policy Project, "Expanding the Vision of High Performance Computing and Communications: Linking America for the Future, Washington, 1991.

[29] Office of Science and Technology Policy, Grand Challenges 1992: High Performance Computing and Communications, January, 1991, page 19.

[30] NEARnet later received a grant to establish a direct connection to the new T-3 backbone.

[31] They have sometimes been referred to as "peer networks" rather than mid-level networks. This reveals the fundamental illogic of the hierarchical model since it makes to no sense to discriminate between the two categories based on the exact geographic scope of the network.

[32] Such is the case in Wyoming. The Wyoming Higher Education Computer Network provides community college users with university addresses and a connection to WestNet.

[33] At the very low end, many commercial bulletin boards are now providing Internet mailing addresses and mail service at flat monthly or quarterly rates, but, again, users without interactive services are not properly on the Internet. Even these rates, $40-80 per year for unlimited e-mail, pale beside the per user costs attributable to university BITNET connections of $5-20 per year for full Internet functionality

[34] A dedicated SLIP port at 9600/14400 baud from Netcom is $160 month plus a setup fee of $750 Local loop costs (dial-up or leased line) are additional.

[35] In fact, NSF's policies were fairly liberal. Commercial traffic was nominally precluded -- but could be approved if it was in support of research and education. NSF almost always approved such proposed uses, but on an "experimental basis." This was not very assuring to potential users, and ANS was able to commit equivalent use of the backbone on an ongoing basis.

[36] See "A Mid-level's Guide to the ANS Agreements," Advanced Network and Services, Inc., August 14, 1991 (2pp.).

[37] CERFnet was and is operated by General Atomics; PSI was a venture capital-funded spinoff of NYSERNET; AlterNet is owned by UUNet Technologies, which was originally set up as a nonprofit corporation.

[38] The implementation of the T-3 backbone has in fact been fraught with technical problems. Although the transition to the T-3 backbone was scheduled to begin in December of 1990, as of this writing the T-1 backbone still carries most of the traffic.

[39] NSF responded by referring to Section 3(a)(4) of the National Science Foundation Act of 1950, as amended, which directs NSF to "foster and support the development and use of computer and other scientific and engineering methods and technologies, primarily for research and education in the sciences and engineering," but concluded: "The AUP may be more restrictive than is legally required, and it is currently being reviewed for possible revision." It is worth recalling that no AUP is imposed on the NSF-subsidized mid-level networks.

[40] P.L. 102-588, Section 217 (42 USC 1862(g)), which states ". . . the Foundation is authorized to foster and support access by the research and education communities to computer networks which may be used substantially for purposes in addition to research and education in the sciences and engineering, if the additional uses will tend to increase the overall capabilities of the networks to support such research and education activities." This change in law helped provide a sound basis for funding Network Access Points that would provide an AUP-free transfer point for exchanging traffic among networks under the future NSFNET architecture.

[41] Program Solicitation: Network Access Point Manager, Routing Arbiter, Regional Network Providers and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program, National Science Foundation, Washington, DC, May 6, 1993.

[42] Meanwhile, the follow-on to FTS 2000 is under study. See, "Post-FTS2000 Acquisition Alternatives White Paper," available on the Internet at post.fts2k.gsa.gov and from the General Services Administration FTS2000 Program Office, Room 6223, 18th and G Sts, NW, Washington, DC 20405.