

Commercialization of the Internet
Summary Report

Status of this Memo

This memo is based on a workshop held by the Science, Technology and Public Policy Program of the John F. Kennedy School of Government, Harvard University, March 1-3, 1990.

This memo provides information for the Internet community. It does not specify any standard. Distribution of this memo is unlimited.

Introduction

"The networks of Stages 2 and 3 will be implemented and operated so that they can become commercialized; industry will then be able to supplant the government in supplying these network services." -- Federal Research Internet Coordinating Committee, Program Plan for the National Research and Education Network, May 23, 1989, pp. 4-5.

"The NREN should be the prototype of a new national information infrastructure which could be available to every home, office and factory. Wherever information is used, from manufacturing to high-definition home video entertainment, and most particularly in education, the country will benefit from deployment of this technology.... The corresponding ease of inter-computer communication will then provide the benefits associated with the NREN to the entire nation, improving the productivity of all information-handling activities. To achieve this end, the deployment of the Stage 3 NREN will include a specific, structured process resulting in transition of the network from a government operation a commercial service." -- Office of Science and Technology Policy, The Federal High Performance Computing Program, September 8, 1989, pp. 32, 35.

"The National Science Foundation shall, in cooperation with the Department of Defense, the Department of Energy, the Department of Commerce, the National Aeronautics and Space Administration, and other appropriate agencies, provide for the establishment of a national multi-gigabit-per-second research and education computer network by 1996, to be known as the National Research and Education Network, which shall:

- (1) link government, industry, and the education

community;

....

(6) be established in a manner which fosters and maintains competition and private sector investment in high speed data networking within the telecommunications industry;

....

(8) be phased out when commercial networks can meet the networking needs of American researchers."

-- S. 1067, 101st Congress, 2nd Session, as marked up April 3, 1990 ["High-Performance Computing Act of 1990"], Title II, Section 201.

Background

This report is based on a workshop held at the John F. Kennedy School of Government, Harvard University March 1-3, 1990, by the Harvard Science, Technology and Public Policy Program. Sponsored by the National Science Foundation and the U.S. Congress Office of Technology Assessment, the workshop was designed to explore the issues involved in the commercialization of the Internet, including the envisioned National Research and Education Network (NREN). Rather than recapitulate the discussion at the workshop, this report attempts to synthesize the issues for the benefit of those not present at the workshop. It is intended for readers familiar with the general landscape of the Internet, the NSFNET, and proposals and plans for the NREN.

At the workshop, Stephen Wolff, Director of the NSF Division of Networking and Communications Research and Infrastructure, distinguished "commercialization" and "privatization" on the basis of his experience developing policy for the NSFNET. He defined commercialization as permitting commercial users and providers to access and use Internet facilities and services and privatization as the elimination of the federal role in providing or subsidizing network services. In principle, privatization could be achieved by shifting the federal subsidy from network providers to users, thus spurring private sector investment in network services. Creation of a market for private vendors would in turn defuse concerns about acceptable use and commercialization.

Commercialization and Privatization

Commercialization. In the past, many companies were connected to the old ARPANET when it was entirely underwritten by the federal government. Now, corporate R&D facilities are already connected to, and are sometimes voting members of, mid-level networks. There are mail connections from the Internet to commercial services such as

MCIMAIL, SprintMail, and Compuserve. DASnet provides a commercial mail gateway to and from the Internet and commercial mail services. UUNET, a nonprofit corporation, markets TCP/IP services (Alternet) with access to the Internet as well as mail services. Performance Systems International (PSI), a startup company which now operates NYSERNET (the New York State regional network, partially funded by NSF) is aggressively marketing Internet-connected TCP/IP services on the East and West Coasts. RLG is selling access to its RLIN database over the Internet directly to end users. Other fee-based services include Clarinet, a private news filtering service, and FAST, a non-profit parts brokering service. However, in all these cases, any use of the NSFNET backbone must, in principle, support the "purpose of the NSFNET."

Under the draft acceptable use policy in effect from 1988 to mid-1990, use of the NSFNET backbone had to support the purpose of "scientific research and other scholarly activities." The interim policy promulgated in June 1990 is the same, except that the purpose of the NSFNET is now "to support research and education in and among academic institutions in the U.S. by access to unique resources and the opportunity for collaborative work." Despite this limitation, use of the NSFNET backbone has been growing at 15-20% per month or more, and there are regular requests for access by commercial services. Even though such services may, directly or indirectly, support the purposes of the NSFNET, they raise prospects of overburdening network resources and unfair competition with private providers of network services (notably the public X.25 packet-switched networks, such as SprintNet and Tymnet).

Privatization. In some respects, the Internet is already substantially privatized. The physical circuits are owned by the private sector, and the logical networks are usually managed and operated by the private sector. The nonprofit regional networks of the NSFNET increasingly contract out routine operations, including network information centers, while retaining control of policy and planning functions. This helps develop expertise, resources, and competition in the private sector and so facilitates the development of similar commercial services.

In the case of NSFNET, the annual federal investment covers only a minor part of the backbone and the regional networks. Although the NSFNET backbone is operated as a cooperative agreement between NSF and Merit, the Michigan higher education network, NSF contributes less than \$3 million of approximately \$10 million in annual costs. The State of Michigan Strategic Fund contributes \$1 million and the balance is covered by contributed services from the subcontractors to Merit, IBM and MCI.

At the regional level, NSF provides approximately 40% of the operating costs of the mid-level networks it funds -- with the remainder covered by membership and connection fees, funding from state governments, and in-kind contributions. This calculation does not include a number of authorized networks (e.g., PREPnet, and, until recently, NEARnet and CERFnet) that receive no NSF funding. However, NSF also funds institutional connections to the NSFNET, which includes payments by the institution to the regional network. Other agencies (DOD, NASA, and DOE) have also funded some connections to NSFNET networks for the benefit of their respective research communities -- and have occasionally funded the networks directly.

Finally, the campus-level networks at academic institutions probably represent a perhaps 7-10 times larger annual investment than the mid-level networks and the backbone together, yet there is no federal funding program at this level. Furthermore, since these local networks must ordinarily be built by the institution rather than leased, there is an additional capitalization cost incurred by the institutions, which, annualized and aggregated, is perhaps another 20-50 times the annual costs of the mid-level and backbone networks. (These figures are the roughest of estimates, intended only for illustration.)

The NSFNET Backbone as a Free Good

Whereas the NSF funding of mid-level networks varies greatly -- from 0% to 75% -- the backbone is available as a free good to the NSF-funded mid-level networks. It is also used free of charge by other authorized networks, including networks not considered part of NSFNET: CSNET, BITNET, UUNET, and PSI, as well as the research networks of other federal agencies. As noted, their use of the backbone is in principle limited to the support of academic research and education.

Through their use of the NSFNET backbone, these networks appear to be enjoying a subsidy from NSF -- and from IBM, MCI, and the State of Michigan. BITNET and some agency networks even use the backbone for their internal traffic. Nonetheless, these other networks generally add value to NSFNET for NSFNET users and regional networks insofar as all networks benefit from access to each other's users and resources.

However, small or startup networks generally bring in fewer network-based resources, so one side may benefit more than the other. To the extent that the mail traffic is predominantly mailing lists (or other information resources) originating on one network, questions of imbalance and implicit subsidy arise. For example, because of the mailing lists available without charge on the Internet, three times as much traffic runs over the mail gateway from the Internet to

MCIMAIL as from MCIMAIL to the Internet. This pattern is reinforced by the sender-pays fee structure of MCIMAIL, which discourages mailing list distribution from within MCIMAIL.

The impact of such imbalances is not clear. For now, the capacity of the NSFNET backbone is staying ahead of demand: It jumped from 56 Kbps to 1.544 Mbps (T-1) in 1988 and will go to 45 Mbps over the next year. But NSF is concerned about a possible recurrence of the congestion which drove users off the NSFNET prior to the 1988 upgrade. Given the tripling of campus-level connections over the past year, continued growth in users at each site, the parade of new resources available over the network, and, especially, the development of high-bandwidth uses, there is reason to fear that demand may again overwhelm capacity.

Offering the NSFNET backbone at no cost to authorized networks both encourages undisciplined use of the backbone and inhibits private investment in backbone networks. It constrains the development of a market for commercial TCP/IP services by diverting an established and rapidly growing user base to a subsidized resource. Charging NSFNET regionals and other mid-level networks for the use of the NSFNET backbone would resolve this problem, but this would impose a substantial cost burden on the mid-level networks, which would in turn have to raise membership and connection fees dramatically. To compensate, the NSF subsidy that now underwrites the backbone could be moved down the distribution chain to the users of the backbone -- i.e., to the regional networks, to the campuses, or even to researchers themselves.

Each option poses unique opportunities and problems. In theory, the further down the chain the subsidy is pushed, the more accountable providers will be to end-user needs. Funding in hands of researchers would make universities more responsive to researchers' networking needs. Funding in the hands of universities would in turn make regional networks more responsive and competitive. And funds for regional networks would spur a general market for backbone services. But the mechanisms for expressing user demand upward through these tiers are imperfect. And, from an administrative standpoint, it is easier for NSF to simply provide one free backbone to all comers -- rather than deal with 25 mid-level networks, or 500 universities, or perhaps tens or hundreds of thousands of individual researchers.

Option: Funding Researchers

It would be possible to earmark funds for network services in agency research grants as a matter of course, so that no new administrative process would be required. But since network costs are presently not usage based, such funding will not readily translate into

identifiable services and may simply end up in local overhead accounts since few institutions allocate out costs of access to the Internet. The use of vouchers rather than cash add-ons might help ensure that federal resources are in fact applied to qualifying wide area network services -- and possibly avoid the imposition of standard institutional overhead on direct funding. However, if vouchers can be sold to other institutions, as economists would advocate in the interests of market efficiency, these advantages may be compromised. Even non-transferable vouchers may create a unique set of accounting problems for both funding agencies and institutional recipients.

A federal subsidy channeled automatically to research grants could substantially limit or segregate the user community. It would tend to divide the academic community by exacerbating obvious divisions between the resource-rich and resource-poor -- between federally funded researchers and other researchers, between scientists and faculty in other disciplines, and between research and education. Within the academic community, there is considerable sentiment for providing basic network services out of institutional overhead to faculty and researchers in all disciplines, at least as long as basic services remain unmetered and relatively low at the institutional level. Of course, special costing and funding may well make sense for high-bandwidth usage-sensitive network services (such as remote imaging) as they become available in the future.

Option: Funding Institutions

Alternatively, funding for external network services, whether in the form of cash or vouchers, could be provided directly to institutions without linking it directly to federal research funding. As it is, institutions may apply for one-time grants to connect to regional networks, and these are awarded based on peer assessment of a number of different factors, not just the quality of the institution's research. But redirecting the subsidy of the backbone could provide regular support at the institutional level in ways that need not involve peer review. For example, annual funding might be tied to the number of PhD candidates within specific disciplines -- or to all degrees awarded in science. Geographic location could be factored in -- as could financial need. This, of course, would amount to an entitlement program, a rarity for NSF. Nonetheless, it would allow institutions to make decisions based on their own needs -- without putting NSF in the position of judging among competing networks, nonprofit and for-profit.

There are, however, questions about what sort of services the earmarked funding or vouchers could be used for. Could they be used to pay the institution's BITNET fee? Or a SprintNet bill? Or to

acquire modems? For information services? And, if so, what sort? Such questions force the funding agency to assume a kind of regulatory in an environment where competing equities, demonstrated need, technological foresight, and politics must be constantly weighed and juggled.

Option: Funding Regional Networks

Shifting the subsidy to the regional networks is appealing in that it appears to be the least radical alternative and would only require allocating funds among some two dozen contenders. Since most of the regional networks are already receiving federal funding, it would be relatively simple to tack on funds for the purchase of backbone services. However, providing additional funding at this level highlights the problem of competition among mid-level networks.

Although most regional networks are to some degree creatures of NSF, funded to ensure the national reach of NSFNET, they do not hold exclusive geographic franchises, and in some areas, there is competition between regionals for members/customers. NSF grants to regional networks, by their very size, have an effect of unleveling the playing field among regionals and distorting competitive strengths and weaknesses.

Alternet and PSI further complicate the picture, since there is no clear basis for NSF or other agencies to discriminate against them. The presence of these privately funded providers (and the possibility of others) raises difficult questions about what network services the government should be funding: What needs is the market now capable of meeting? And where will it continue to fail?

Experience with regulation of the voice network shows that it is inefficient to subsidize local residential service for everybody. If one is concerned about people dropping off the voice network -- or institutions not getting on the Internet -- the answer is to identify and subsidize those who really need help. The market-driven suppliers of TCP/IP-based Internet connectivity are naturally going after those markets which can be wired at a low cost per institution, i.e., large metropolitan areas, especially those with a high concentration of R&D facilities, such as Boston, San Francisco, and Washington, DC. In the voice environment, this kind of targeted marketing by unregulated companies is widely recognized as cream-skimming.

Like fully regulated voice common carriers (i.e., the local exchange carriers), the non-profit NSF-funded regional networks are expected to serve all institutions within a large geographic area. In areas with few R&D facilities, this will normally result in a

disproportionately large investment in leased lines. Either remote institutions must pay for the leased line to the nearest network point of presence -- or the network must include the leased line as part of common costs. If the regional network assumes such costs, it will not be price-competitive with other more compact networks.

Accordingly, a subsidy redirected to the regional networks could be keyed to the density of the network. This might be calculated by number of circuit miles per member institution or some form of aggregate institutional size, figured for either the network as a whole or for a defined subregion. This subsidy could be available to both for-profit and non-profit networks, but only certain non-profit networks would meet the density requirement, presumably those most in need of help.

Increasing the Value of the Connection

The principal advantage in underwriting the backbone is that it provides a evenhanded, universal benefit that does not involve NSF in choosing among competing networks. By increasing the value of belonging to a regional network, the backbone offers all attached networks a continuing annual subsidy commensurate with their size.

Increased value can also derived from access to complementary resources -- supercomputer cycles, databases, electronic newsletters, special instruments, etc. -- over the network. Like direct funding of backbone, funding these resources would induce more institutions to join regional networks and to upgrade their connections. For example, where a database already exists, mounting it on the network can be a very cost-effective investment, increasing the value of the network as well as directly benefiting the users of the database.

Commercial information services (e.g., Dialog, Orbit, Lexis) may serve this function well since they represents resources already available without any public investment. Marketing commercial services to universities over the Internet is permissible in that it supports academic research and education (although the guidelines state that such commercial uses "should be reviewed on a case-by-case basis" by NSF).

But to date there has been remarkably little use of the regional networks, let alone the NSFNET backbone, to deliver commercial information services. In part, this is because the commercial services are unaware of the opportunities or unsure how to market in this environment and are concerned about losing control of their product. It is also due to uneasiness within the regional networks about usage policies and reluctance to compete directly with public packet-switched networks. However, for weak regional networks, it

may be necessary to involve commercial services in order to attract and hold sufficient membership -- at least if NSF subsidies are withdrawn. Without a critical mass of users, commercialization may need to precede privatization.

Impact of Removing NSF Subsidy from the Backbone

Any shift to a less direct form of subsidy may cause some dislocation and distress at the regional network level -- until the benefits begin to be felt. No regional network has yet folded, and no institution has permanently dropped its connection to a regional network as a consequence of higher prices, but concerns about the viability of some regionals would suggest that any withdrawal of subsidy proceed in phases.

Moreover, as the NSF subsidy vanishes, the operation of the backbone becomes a private concern of Merit, the Michigan Strategic Fund, IBM, and MCI. While Merit and the Michigan Strategic Fund are more or less public enterprises within the state, they are essentially private entrepreneurs in the national operation of a backbone network. Without NSF's imprimatur and the leveraging federal funds, the remaining parties are much less likely to treat the backbone as a charity offering and may well look to recovering costs and using revenues to expand service.

The backbone operation could conceivably become either a nonprofit or for-profit utility. While nonprofit status might be more appealing to the academic networking community now served by the backbone, it is not readily apparent how a broadly representative nonprofit corporation, or even a cooperative, could be constituted in a form its many heterogeneous users would embrace. A non-profit organization may also have difficulty financing rapid expansion of services. At the same time, the fact that it will compete with private suppliers may preclude recognition as a tax-exempt organization -- and so its ability to reinvest retained earnings.

Operation of the backbone on a for-profit basis would attract private investment and could be conducted with relative efficiency. However, given the dominant position of the backbone, a for-profit operation could conceivably get entangled in complex antitrust, regulatory, and political struggles. A nonprofit organization is not immune from such risks, but to the extent its users are represented in policy-making, tensions are more likely to get expressed and resolved internally.

The status of backbone or regional networks within the Internet is entirely separate from the question of whether network services are metered and charged on a usage basis. Confusion in this regard stems

from the fact that the low-speed public data networks (SprintNet, TymNet), which are sometimes seen as competitive to Internet services, do bill on a connect-time basis. However, these commercial services use X.25 connection-based packet-switching -- rather than the connectionless (datagram) TCP/IP packet-switching used on the Internet. Internet services could conceivably be billed on per-packet basis, but the accounting overhead would be high and packets do not contain information about individual users. At bottom, this is a marketing issue, and there is no evidence of any market for metered services -- except possibly among very small users. The private suppliers, Alternet and PSI, both sell "pipes" not packets.

Privatization by Function

As an alternative approach to encouraging privatization, Dr. Wolff suggested barring mature services such as electronic mail from the subsidized network. In particular, NSF could bar the mail and news protocols, SMTP and NNTP, from the backbone and thereby encourage private providers to offer a national mail backbone connecting the regional networks. Implementation would not be trivial, but it would arguably help move the academic and research community toward the improved functionality of X.400 standards. It would also reduce traffic over the backbone by about 30% -- although given continued growth in traffic, this would only buy two months of time.

If mail were moved off the regional networks as well as off the NSFNET backbone, this would relieve the more critical congestion problem within certain regions. But logistically, it would be more complicated since it would require diverting mail at perhaps a thousand institutional nodes rather than at one or two dozen regional nodes. Politically, it would be difficult because NSF has traditionally recognized the autonomy of the regional networks it has funded, and the networks have been free to adopt their own usage guidelines. And it would hurt the regional networks financially, especially the marginal networks most in need of NSF subsidies. Economies of scale are critical at the regional level, and the loss of mail would cause the networks to lose present and potential members.

The National Research and Education Network

The initiative for a National Research and Education Network (NREN) raises a broader set of policy issues because of the potentially much larger set of users and diverse expectations concerning the scope and purpose of the NREN. The decision to restyle what was originally described as a National Research Network to include education was an important political and strategic step. However, this move to a broader purpose and constituency has made it all the more difficult

to limit the community of potential users -- and, by extension, the market for commercial services. At the regional, and especially the state level, public networking initiatives may already encompass economic development, education at all levels, medical and public health services, and public libraries.

The high bandwidth envisioned for the NREN suggests a growing distance between resource-intensive high-end uses and wide use of low-bandwidth services at low fixed prices. The different demands placed on network resources by different kinds of services will likely lead to more sophisticated pricing structures, including usage-based pricing for production-quality high-bandwidth services. The need to relate such prices to costs incurred will in turn facilitate comparison and interconnection with services provided by commercial vendors. This will happen first within and among metropolitan areas where diverse user needs, such as videoconferencing and medical imaging, combine to support the development of such services.

As shown in Figures 1. and 2., the broadening of scope corresponds to a similar generalization of structure. The path begins with mission-specific research activity organized within a single computer. It ends with the development of a national or international infrastructure: a ubiquitous, orderly communications system that reflects and addresses all social needs and market demand, without being subject to artificial limitations on purpose or connection. There is naturally tension between retaining the benefits of specialization and exclusivity and seeking the benefits of resource-sharing and economies of scale and scope. But the development and growth of distributed computing and network technologies encourage fundamental structures to multiply and evolve as components of a generalized, heterogeneous infrastructure. And the vision driving the NREN is the aggregation and maturing of a seamless market for specialized information and computing resources in a common, negotiable environment. These resources have costs which are far greater than the NREN. But the NREN can minimize the costs of access and spread the costs of creation across the widest universe of users.

Figure 1. Generalization of Purpose:

Discipline-Specific Research	CSNET, HEPnet, MFEEnet
General Research	early NSFNET, "NRN"
Research and Education	BITNET, present NSFNET, early "NREN"
Quasi-Public	many regional networks, "NREN"
National Infrastructure	"commercialized NREN"

Figure 2. Generalization of Structure:

Computer	time-sharing hosts
Network	early ARPANET
Internetwork	ESNET, NSFNET (tiered)
Multiple Internetworks	present Internet
Infrastructure	"NREN"

Workshop Participants

Rick Adams, UUNET
 Eric Aupperle, Merit
 Stanley Besen, RAND Corporation
 Lewis Branscomb, Harvard University
 Yale Braunstein, University of California, Berkeley
 Charles Brownstein, National Science Foundation
 Deborah Estrin, University of Southern California
 David Farber, University of Pennsylvania
 Darleen Fisher, National Science Foundation
 Thomas Fletcher, Harvard University
 Kenneth Flamm, Brookings Institution
 Lisa Heinz, U.S. Congress Office of Technology Assessment
 Fred Howlett, AT&T
 Brian Kahin, Harvard University
 Robert Kahn, Corporation for National Research Initiatives
 Kenneth King, EDUCOM

Kenneth Klingenstein, University of Colorado
Joel Maloff, CICNet
Bruce McConnell, Office of Management and Budget
Jerry Mechling, Harvard University
James Michalko, Research Libraries Group
Elizabeth Miller, U.S. Congress Office of Technology Assessment
Eli Noam, New York State Public Service Commission
Eric Nussbaum, Bellcore
Peter O'Neil, Digital Equipment Corporation
Robert Powers, MCI
Charla Rath, National Telecommunications and Information
Administration, Department of Commerce
Ira Richer, Defense Advanced Research Projects Agency
William Schrader, Performance Systems International
Howard Webber, Digital Equipment Corporation
Allan Weis, IBM
Stephen Wolff, National Science Foundation

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Brian Kahin
Director, Information Infrastructure Project
Science, Technology & Public Program
John F. Kennedy School of Government
Harvard University

Phone: 617-495-8903

EMail: kahin@hulaw.harvard.edu