Application of Testbed Research 5

The networking component of the High Performance Computing and Communications (HPCC) Program funds both research on gigabit technology and the deployment of this technology in the National Research and Education Network (NREN). One of the NREN's roles is to provide additional experience with advanced network technologies before they are deployed more widely in the national information infrastructure. However, the testbed research will also be applied directly to other networks, such as the common carriers' public switched network, without intermediate deployment in the NREN.

APPLICATION TO THE NREN

There is no overall NREN development plan; however, the National Science Foundation (NSF) is to coordinate the evolution of the Federal agency networks that are the core of the NREN.[1] During 1992, NSF, the Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA) announced plans for the future development of their networks.[2] Based on these plans, the next-generation agency networks will likely be similar to the testbed networks, with an emphasis on Synchronous Optical Network (SONET) fiberoptic transmission and fast packet switching. These broadband technologies are

The testbed research can be applied to networks other than the NREN.

[1] Office of Science and Technology Policy, "Grand Challenges 1993: High Performance Computing and Communications," p. 33.

[2] National Science Foundation, "Public Draft: Network Access Point Manager/ Routing Authority and Very High Speed Backbone Network Services provider for NSFNET and the NREN Program," June 12, 1992; James F. Leighton, Manager of Networking and Engineering, National Energy Research Supercomputer Center, Lawrence Livermore National Laboratory, "ESnet Fast-Packet Services Requirements Specification Document," Feb. 20, 1992.

being studied in a large number of research programs, but the testbed research is unique in its emphasis on building wide area gigabit networks and testing them with applications.

Agency Plans

The Federal agencies will not build their own ''private' networks, but will obtain services from a network service provider. In effect, NSF, DOE, and NASA will act as early, large customers for advanced services. While industry has developed the switches and transmission equipment required for advanced network services, the agency backbones will be one of the first opportunities to integrate these components into a system that provides services to real users. Users of the agency backbones are knowledgeable about networking and will assist in integrating new network services with computers and applications to create useful systems.

The agency backbone services could be provided by a number of different organizations—the carriers, computer companies, or some of the emerging providers of commercial Internet services-or consortia. Provision of services for agency backbone networks provides valuable experience that the network operator may be able to translate into earlier availability of advanced services on a commercial basis.[3] For prospective players, the decision to participate in the provision of services to the agency networks weighs the experience gained and long-term strategic considerations against the cost of providing the service, which is greater than the money available from the Federal agencies.

To help stimulate market interest, DOE and NASA had originally decided to combine their NREN-related programs. A single supplier would have provided network services to both agencies, connecting sites such as DOE's Los Alamos National Laboratory or NASA's Ames Research Center. However, the General Accounting Office (GAO) overturned DOE's choice of contractor in March, 1993 (see ch. 1, p. 7). The steps that the agencies will take in response to this decision were still unclear at the time of publication, but it is possible that DOE and NASA will now decide to proceed separately. The procurement process has been significantly delayed, and will likely not be completed before the fourth quarter of 1993. Before the GAO decision, NASA and DOE had planned to begin connecting sites to the new network in mid-1993.

NSF issued a draft solicitation for its next-generation network in mid-1992. NSF plans to publish a final version of the solicitation and award a cooperative agreement during 1993. The new network is scheduled to begin operation in mid-1994. NSF's plans for the evolution of its network have greater implication for the evolution of the NREN and the Internet than do those of DOE and NASA. The current NSFNET backbone carries much more traffic than the other agency backbones [4] and serves a broader range of users. However, many of those users will not be able to use the next-generation backbone.

The new NSF network's Primary purpose will be to connect the NSF supercomputer centers, enabling advanced distributed supercomputing applications. By contrast, today's NSFNET backbone is a' 'general-purpose' network that carries all types of research and education traffic. NSF envisions that in the future routine research and education traffic will be handled by commercial providers, not by the agency-operated backbone. There are a number of emerging commercial providers, and the network technology required to support routine traffic is sufficiently stable and reliable. This strategy would also free capacity on the backbone for experimental applications.

The transition to the new environment resulting from the changed role of the NSFNET

---

[3] Ann H. Lindstrom, "Sprint Blasts Ahead With ATM Deployment, " *Telephony,* vol. 223, No. 8, Aug. 24, 1992.

[4] Stephen S. Wolff, Director, Division of Networking and Communications Research and Infrastructure, National Science Foundation testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, Serial No. 120, p. 155.

backbone will require careful management to ensure stability. NSF's plan will affect significantly the existing three-level hierarchy of the NSFNET. The regional networks were designed to provide connections to sites on the current backbone, which in turn provides inter-regional connectivity. Under NSF's new plan, the backbone will serve many fewer sites and will no longer play the same central role in research and education networking. The regional networks will have to make new arrangements for interconnections and will be operating in a more competitive environment.

Agency Backbone Technology

The collaborative nature of the testbeds makes it more likely that the network technologies developed by industry will be suitable for operation in the agency backbone environment. The testbeds are emphasizing the technologies' use with the Internet protocols used by the agency networks, and are studying the interaction between fast packet networks and supercomputer network standards and applications. In addition, they emphasize the gigabit bandwidths required to support the Grand Challenge applications that are a key component of the overall HPCC program. The involvement of the carriers in the research program may also lead the carriers toward a more active role in providing NREN services.

While the plans for the evolution of the agency backbones are consistent with the target established by the testbeds, the agency networks will initially operate at lower bandwidths than the testbeds. The agency backbones will incorporate more of the technology from the testbed research as they evolve over time to meet the goal of the gigabit NREN. However, some issues cannot be addressed by the testbeds, or may be discovered only as the agency networks are deployed. Many of these issues are related to the more complex topologies (greater number of sites), larger number of users, and more diverse sources of traffic that will be present on the production networks.

TRANSMISSION TECHNOLOGY

The agencies envision the use of SONET equipment similar to that used in the testbeds, and have indicated that they hope to use 155 Mb/s SONET equipment in 1994 and then upgrade over time to 622 Mb/s (the next transmission rate in the SONET family) by 1996, the High Performance Computing Act's target year for the use of gigabit links. The 622 Mb/s rate, less than a full gigabit per second, is sometimes referredtoasa‘‘goverment gigabit."[5]

The rate at which the agency backbones will evolve depends on the timely deployment of a high-bandwidth SONET transmission infrastructure by the carriers. While much of the carriers' existing network uses fiber, SONET transmission equipment is required in order to support computer networking above the current T3 rates-it allows the fiber to be configured to carry high bandwidth channels. However, this equipment is extremely costly at this time and the carriers' deployment schedules have been slipping from earlier estimates.

The testbed networks will have also provided experience with the connection of supercomputers to high speed networks. ‘‘High end" users will require fiber links connecting their sites to the NREN. Only fiber is able to carry the large amounts of data needed for supercomputer-based applications. The testbeds are one of the first large-scale deployments of SONET to end-users, and considerable work has been done on interface devices to connect supercomputers and high-speed local area networks to fast packet switched networks. However, widespread use of high-speed networks will depend in part on the degree to which computer companies design their workstations to be fully integrated into a high-speed network. Today, bottlenecks encountered in mov-

[5] Carpenter et al., "Where Are We With Gigabits?" *IEEE Network, vol. 6, No, 2,* March 1992, p. 10.

ing data from the network into the computer's memory, where it can be used by the applications software, can limit the performance of the overall system.

SWITCHES

The next-generation backbone networks will use fast packet switching technology similar to that used in the testbeds. Initially, the switches will not be as sophisticated, because of the lower link bandwidths. The network operator's choice of switching technology, from among those being investigated in the testbeds and elsewhere, depends in part on long-term strategic considerations. If a carrier were to provide services for an agency network, it would probably use Asynchronous Transfer Mode (ATM) switches. ATM has been chosen as the foundation for the future development of carrier networks, and the provision of services for the agency backbone would provide an opportunity to gain experience with its use. Other providers might also choose to use ATM switches, or strategic considerations may lead to the choice of an alternate switching technology.

The DOE Reqest for Proposals issued in early 1992 specified the use of fast packet "cell relay' technology. "Cell relay" is a term used to describe both ATM and Switched Multimegabit Data Service (SMDS), a data communications service developed by the telephone companies. In the summer of 1992, DOE and NASA selected a contractor that proposed to use ATM. This DOE/NASA program would have been the frost large-scale deployment of ATM. One of the goals of DOE and NASA is to encourage the development of commercial services by evaluating and demonstrating emerging technologies such as ATM. The agencies' effectiveness in performing this function may be reduced by any further delays resulting from GAO's decision overturning their choice of contractor.

The National Science Foundation's draft solicitation describing the evolution of its backbone network did not specify a particular type of switch.[6] NSF will allow prospective bidders to propose their choice of switching technology. The most likely option that would be proposed would be an ATM-based approach. Another type of fast packet technology, such as the PTM approach developed by one of the participants in the Aurora testbed, might also be used. The approach of 'overlaying' an Internet network on a network that uses fast packet technology is not unique to ATM. However, ATM has broad support from industry standards committees.

OTHER NREN NETWORKS

The regional networks and other commercial providers of Internet services may also carry NREN traffic. Operators of these networks are faced with the same technology choices as those for the backbone networks. However, because many of these networks will require lower bandwidths than the backbones, they may continue to use ''router-based" networks or use new "pre-broadband" services being offered by the carriers and other service providers. Two examples of these pre-broadband services are Frame Relay and SMDS.[7] These are packet switching services that can also be used to carry Internet traffic (see ch. 2, p. 34). Because the Internet protocols are able to hide differences in network technology from the users of the network, the NREN's networks can be based on a variety of different technologies.

Campus networks and other networks based primarily on local area networks will also become more capable. Local area network research is not currently a focus of the testbeds, although the interconnection of local and wide area networks

6 National Science Foundation, ''Request for Public Comment: Solicitation Concept NSFNET Backbone Network Services, " June 12, 1992.

7 Terry Sweeney, ''TCP/IP-SMDS Trial Completed, " *CommunicationsWeek,* Aug. 17, 1992, p. 25.

is being studied. One of the most important trends in local area network design is that there is a growing amount of support for ATM-based local area networks and products for ATM local area networks are beginning to appear. Other kinds of high-bandwidth local area network standards are also being studied by standards committees.

Applications

Because of the emphasis on gigabit applications, the testbed applications research is primarily applicable to high-end users of the NREN. The testbeds have been one of a small number of research programs to address supercomputer-related networking issues. These applications are, in general, of little concern to industry and would receive less attention without the testbeds. The testbeds' gigabit applications research will have important impacts on the overall HPCC initiative. Distributed supercomputing maybe an important tool for bringing more processing power to bear on the Grand Challenge problems. In addition, the Grand Challenge teams will be scattered about the country and could use networks to support collaboration. The sizes of the data sets used in Grand Challenge problems will be very large, requiring high-bandwidth networks to move them from place to place within a reasonable period of time.

High-speed network support of supercomputing is important to the missions of the NSF supercomputer centers and the Federal laboratories. Ied by testbed participants, the NSF supercomputer centers have proposed a concept that would make use of the distributed supercomputing ideas investigated by the testbeds.[8] They envision a ''metacenter' '—the use of the new high-speed backbone to integrate the computational and intellectual resources of the supercomputer centers.[9] In effect, it would be possible for the four supercomputer centers to act as a single center, distributing a computation among several machines as the computation required.

High-end users of the agency backbones are only part of the user community addressed by the NREN program. Few users will have access to a full gigabit/second of bandwidth, and the supercomputer applications studied by the testbeds are by definition highly specialized. For most users the primary result of improving network capability will be better performance with existing applications and the wider use of video and image-based communications. Because these capabilities may have considerable significance in commercial applications, much work is being done on these types of applications by industry. Some types of applications development, however, may require added support. Legislation introduced in the 103rd Congress (S.4 and H.R. 1757) seeks to expand support for applications development in a variety of education, medicine manufacturing, and library settings.

Internetworking

The NREN is closely linked to the evolution of the Internet protocols that enable the thousands of independently operated networks that make up the Internet to exchange traffic. The testbeds are providing an opportunity to investigate the use of the Internet protocols in fast packet switched networks. The collaborative nature of the testbed program may be encouraging the Internet community to influence the ATM standards process to better support Internet protocols. In addition, the testbeds are investigating the behavior of the Internet protocols at high speeds, and comparing them to some new concepts in protocol design. In the past few years, a number of protocols have been proposed that may perform better in high-speed networks and are better suited to the new fiber-based, fast packet switched networks. For example, today's Internet protocols are designed

[8] John Markoff, "A Crucial Linkup in the US Data Highway," *New York Times*, Sept. 30, 1992, p. D8.

[9] Carolyn Duffy Marson, "NSF Pursues Computing Without Walls,' *Federal Computer Week*, vol. 6, No. 35, Nov. 30, 1992, p. 1.

Figure 5-l-Growth In NSFNET Networks

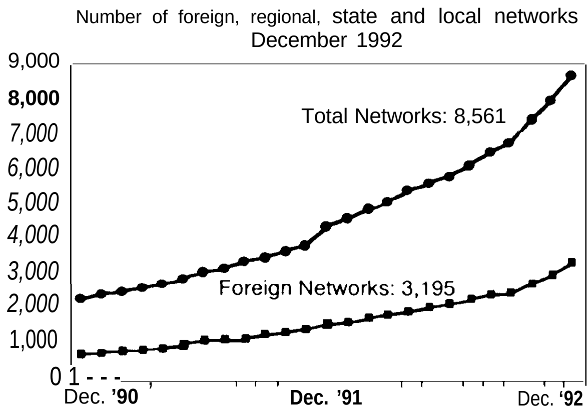

SOURCE: Merit, Inc., 1992.

to handle the types of transmission errors that occur with poor-quality copper lines, but rarely occur with new fiber-based transmission systems.

Other issues related to the evolution of the Internet protocols are not being studied by the testbeds. The main issue confronting the Internet community today is the growing size and complexity of the network—not increases in bandwidth. The growth in the number of users and networks that make up the Internet is putting pressure on current ''routing" technology (figure 5-l). Routing is the process by which a path from one computer to another through a series of intermediate networks is determined. Calculating these paths using current algorithms demands a considerable amount of processing power; the problem is getting worse as the Internet continues to grow and become more complex. Routing issues have not been studied by the testbeds, which only connect a few sites.

Work on issues related to managing the growth of the Internet is being done primarily within the Internet community's technical organizations, such as the Internet Activities Board (IAB) and the Internet Engineering Task Force (IETF). The IETF consists of a number of working groups, one of which addresses routing issues. Currently, within the technical community there are many different proposals; some only address immediate problems, while others attempt to solve the problems in a way that will be satisfactory for a number of years. Besides addressing issues related to growth, some of the new routing algorithms may also take into account the growing diversity of service providers and network capabilities. Routing and management problems associated with the growing Internet are a major research area that requires more study .10

NSF's plan for the evolution of its network as part of the NREN program is linked closely to changes in routing technology. Today, the NSFNET backbone operator plays an especially important role in determining routes for research and education networks. As the Internet becomes more commercialized, however, it becomes less appropriate for NSF to be responsible for this aspect of its operation. NSF envisions reducing the reliance of Internet networks on the NSFNET backbone's operator for routing information.[11] NSF has proposed that the routing function be handled by a separate organization, the "routing authority,' not by the operator of NSF's network. NSF's plan also calls for the creation of a number of Network Attachment Points (NAPs), where commercial networks and agency networks could obtain routing information and interconnect with each other (see box 5-A).

APPLICATION TO OTHER NETWORKS

The testbed program will also impact the evolution of the national information infrastructure more directly, without the intermediate stage of deployment in the NREN. This national information infrastructure includes the larger U.S. Internet-the NREN program targets only one part of the U.S. segment of the Internet (see

10 ARPA is supporting research in these areas, such as through its DARTnet program.

11 Robert Aiken et al., National Science Foundation, ''NSF Implementation Plan for Interagency Interim NREN,'' May 1992.

figure 5-2).[12] It also includes a wide array of other services and technologies to be offered by the carriers, cable television companies, computer hardware and software companies, information service providers and others.

Figure 5-2—Relationship Between NREN, Internet, and National Information Infrastructure

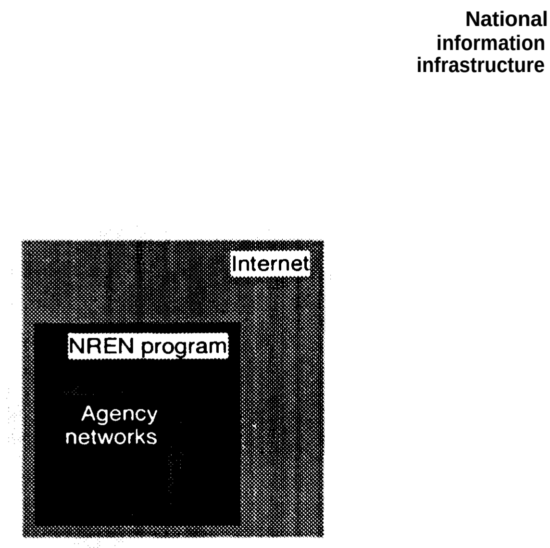

SOURCE: Office of Technology Assessment, 1993.

Application to the Internet

The Internet is increasingly expanding to serve communities other than the core research and education community that is the focus of the NR.EN program. The regional networks and new commercial providers now carry business traffic. The trends towards broader use of the Internet and growing numbers of users seem likely to continue. These will be driven in part by the advances in switching and transmission technology described in chapter 3. They depend to a greater extent on addressing the security concerns of commercial customers, the degree that use of Internet applications can be simplified, and the deployment of advanced digital local loop technologies.

It is possible that the switches and fiber optic links deployed by providers of agency backbone services will also carry commercial traffic. Some of the capacity would be used for the agency backbone network and some would be used to provide services to commercial customers. The Federal agency backbone would be the network's most important customer, acting as a catalyst for the deployment of the required switches and transmission equipment, while commercial customers would help to recover that portion of the costs of operating the network not covered by the Federal agencies' funding.

Other Services

The network technology studied by the testbeds is equally applicable to services other than Internet services. The research will also be applied directly to private networks, the common carriers' public switched network, and possibly cable television networks.[13] This is because the network technology used in the testbeds reflects near-term industry planning. While the testbeds have emphasized higher bandwidths and more specialized applications than have immediate commercial importance, the basic design of the testbed networks-such as the use of fast packet switching and SONET-reflects ideas that figure prominently in industry plans.

The carriers and other network operators could use the new advanced technologies to provide Internet services, or an array of other voice, video, and data communications services. Switch and transmission technologies, though advancing at different rates, are making substantial progress. Because of their commercial importance, fast packet and fiber optic technologies are being

12 For a discussion of the relationship of the Internet to the NREN, see chapter 2, p. 31.

13 For a description of a cable provider's plans to use ATM, see Richard Karpinski, ''Time Warner to LECs: Here We Come,' *Telephony*, vol. 224, No. 5, Feb. 1, 1993, p. 8. IBM is testing its plaNET (PTM) fast packet switch in a cable network in Toronto, Canada. See ' 'IBM's 1.2 Gigabit/sec. Networking Scheme, " *Cable-Telco Report*, August 1992, p. 9.

Box 5-A-NSFNET Backbone Recompetition

The National Science Foundation's plans for the future development of its backbone network have attracted considerable scrutiny. The NSFNET backbone plays an especially important role in the Internet and in the National Research and Education Network program. Currently, NSF has a cooperative agreement with Merit Network, a not-for-profit organization of nine Michigan universities. However, Merit does not operate the NSFNET backbone "in-house." A second organization, Advanced Network& Services (ANS), operates the network-Merit obtains the services for the NSFNET backbone from ANS.

The cooperative agreement with Merit for the NSFNET backbone was announced in November 1987, and covered the 5-year period to November of 1992. Merit's proposal was submitted in partnership with IBM and MCI. The relationship between Merit and its partners changed in September of 1990, when Merit, IBM, and MCI announced the formation of ANS, as a not-for-profit corporation. ANS received capital from MCI and IBM at its formation, and IBM and MCI provide switches, transmission capacity, and other services to ANS. Overtime, more of the responsibilities for the NSFNET backbone have been shifted from Merit to its subcontractor, ANS.

Over the life of the 5-year cooperative agreement with Merit, there have been three important changes in the Internet First, Merit and ANS have Increased the NSFNET backbone's link bandwidth from 56 kb/s, to 1.5 Mb/s (T1), to 45 Mb/s (T3). Second, the Internet has become a much more important part of the U.S. information infrastructure--the amount of traffic and the number of users has grown rapidly. Finally, the past 2 years have seen the emergence of commercial Internet service providers. In particular, ANS has created a for-profit subsidiary. While the T1 network was used only by NSFNET backbone traffic, the T3 network operated by ANS is shared by the NSFNET backbone and ANS's commercial customers.

The relationship between NSF, Merit, ANS, and other commercial providers was the subject of hearings before the House Subcommittee on Science in March of 1992. Concern was expressed by some witnesses that the current arrangement benefited ANS unduly, and had not been foreseen by the 1987 cooperative agreement with Merit. Other witnesses pointed to the success of the NSFNET backbone, the growth in the number of users, and the value of the equipment and services contributed by Merit and its partners.

Recompetition

In preparation for the expiration of the cooperative agreement with Merit in November of 1992, NSF studied a number of options for the future development of the NSFNET backbone. In studying these options, NSF had to take into account several factors that did not apply in 1987. One factor was the emergence of commercial providers. Any new plan for the backbone could not favor the incumbent, ANS, and would have to provide equal opportunity for all firms wishing to provide services to the NSFNET backbone. A second factor was the need for stability. The Internet is now an essential infrastructure for many more users than in 1987, and stability would have to be ensured during the transition to any new arrangement. Finally, NSF had to take into account the NSFNET backbone's central role in the NREN program.

One option studied by NSF was to discontinue direct funding o fa backbone network. Instead, NSF could fund the regional networks and allow them to choose among commercial providers of interconnections, encouraging further development of the commercial networks. According to testimony of the director of the NSF division responsible for NSFNET, this plan was opposed by the regional networks and by other Federal agencies, in part because of concerns about stability during the transition to this environment.

As a result, NSF decided that it would continue to operate a backbone network. NSF's timetable called for extending the arrangement with ANS for up to 18 months beyond November 1992, to the middle of 1994. This eighteen month period was intended to allow time to 1) select the provider of the next-generation NSFNET backbone, and 2) install the required links and switches. Originally, NSF planned to make the awards in the middle of 1993, allowing a year for the transition to the new network,

The NSFNET Solicitation Concept

The Project Development Plan for the continued provision of NSFNHET backbone services after the expiration of the agreement with Merit was published by NSF in November of 1991. This development plan stated the requirements for stability, fair competition, and support of NREN objectives. The Development Plan also presented the concept of splitting the current NSFNET backbone provider's tasks into two parts, and awarding each part to separate organizations.

NSF published a more detailed version of this plan in June of 1992 and requested public comments. According to the plan, one of the two awards would be for the provision of very high speed Backbone Network Services (vBNS). The vBNS provider would operate the i inks and switches and be responsible for moving packets through the NSFNET backbone. Among other requirements, the vBNS provider would establish a network that would operate at 155 Mbps or higher and would "provide for real-time multimedia services, including multicasting and video teleconferencing." NSF did not specify a switching or transmission technology; however, the reference to 155 Mbps implies the use of SONET transmission equipment

The second award would be for the Routing Authority (RA). The routing authority would be responsible for the routing functions that had previously been performed by the backbone operator. The RA would also operate Network Access Points (NAPs), which would facilitate the connection of other networks to the vBNS and to each other. These could be other Federal networks, or commercial networks. The routing information required in order to facilitate the coordination of these networks would be stored in a database accessible at the NAPs. A total of about $10 million annually would be available for the two awards.

Changes to the Draft Solicitation

The public comments received by NSF in response to the draft proposal reflect the degree to which NSF's plans affect more than just the NSFNET backbone. NSF's proposed NAP/RA structure could best be characterized as an "architecture" for the NREN and the internet, with significant implications for the larger information infrastructure. As such, the NSFS plans affected users, interexchange and local exchange carriers, regional networks and other current and prospective providers of Internet services, and other federally supported networks.

As of May, 1993, a revised version of the NSF solicitation had not been released. However, in December 1992, NSF outlined its intention to change its original plan in a number of ways. While the basic vBNS/NAP/RA structure was maintained, NSF indicated that it would make three awards, not two. The NAPs would not have to be operated by the Routing Authority, as had been specified in the draft solicitation, but could be operated by a separate organization.

More importantly, NSF announced that the new backbone would be used primarily to connect the NSF supercomputer centers. The draft solicitation had indicated t hat the new network would continue to be a "general purpose" backbone, serving a large number of sites and carrying both routine and high-end traffic. By Iimiting the scope of the backbone, NSF's new approach would require more routine services to be obtained from commercial providers.

SOURCES: Robert Aikan et al., "NSF Implementation Plan for Interagency Interim NREN," May 1992; National Science Foundation, "Project Development Plan: Continuation and Enhancement of NSFNET Backbone Services," November 1991; National Science Foundation, "Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program: Program Solicitation," June 1992; Ellen Hoffman, "NSFNET Backbone Service Restructured," Link Letter, vol. 5, No. 3, November 1992, p. 1; Douglas E. van Houweling, Merit Network Inc., testimony at hearings before the House Subcommittee on Science, Mar. 12,1992, pp. 36-41; Stephen S. Wolff, Assistant Director, Directorate for Computer and Information Science and Englneerlng, National Science Foundation, testimony at headings before the House Subcommittee on Science, Mar. 12,1992, pp. 133-136, pp. 148-156; William L. Schrader, President and CEO, Performance Systems International, Inc., testimony at hearings before the House Subcommittee on Science, Mar. 12, 1992, pp. 87-98; Ellen Messmer, "NSF Changes Course On Its Internet Plan," Network World, vol. 9, No. 51, Dec. 21, 1992, p. 1; Office of Inspector General, National Science Foundation, "Review of NSFNET," Mar. 23, 1993.

studied by a large number of research programs in addition to the testbeds. The issues affecting the deployment of these technologies in commercial settings are mainly concerned with trading the costs associated with the existing infrastructure against the potential of future markets for the new technologies. Regulatory and economic factors affecting the pace of deployment are beyond the scope of this background paper.

The involvement of the carriers in the testbeds was an important result of the visibility afforded by the HPCC program and the Corporation for National Research Initiatives' organization. All three major interexchange carriers and most of the Regional Bell Operating Companies are involved. The focus on ATM-related issues serves to provide experience with the construction of these networks and demonstrate their feasibility on a significant scale. Despite the carriers' stated commitment to ATM, the degree to which the transition to ATM represents a true paradigm shift for the telecommunications industry should not be underestimated. The testbeds will have served to help advance the carriers beyond the stage of standards-setting, component development, and small-scale experiments. There are many who believe that a nationwide gigabit network is not possible without basing it on the ongoing investments of the carriers.[14]

The testbeds may also be helping to provide input to the ATM standards process. Currently, there is some concern in the telecommunications industry that elements of ATM are being standardized before there is sufficient understanding of the tradeoffs. In particular, there is uncertainty about the best way to control the traffic in ATM networks, a key component in the use of ATM to support integrated services. The testbeds will provide experience with real traffic, due to the involvement of applications researchers. The academic researchers are also contributing to the solution of these problems; while algorithms for the control of packet networks area longstanding topic of theoretical research, the testbeds may serve to focus the work of academic researchers on topics of concern to industry to a greater extent.

14 Vinton G. Cerf, "Another Reading of the NREN Legislation," *Telecommunications*, vol. 25, No. 11, November 1991, p. 29.