




**NSFNET: A Partnership for High-Speed  
Networking  
Final Report  
1987-1995**

# Table of Contents

Introduction .....	4
I. Coming Together .....	5
Merit and the Community to Serve: Research & Education.....	5
The Role of Industry: IBM and MCI .....	7
The Role of Government: The National Science Foundation .....	11
The Vision: Networking the Research and Education Community .....	15
The Solicitation.....	17
Drafting the Proposal and Winning the Award.....	18
II. The Partnership.....	20
Building a National Backbone Service .....	20
Transition to T-3: Challenges and Opportunities.....	27
Milestones in Technology.....	34
The NSFNET Phenomenon .....	35
Education and Outreach: To the Academic Community and Beyond .....	36
Realizing the Vision: Lessons and Insights from the Project Team .....	38
III. Transition to the Future	
A New Architecture for the NSFNET.....	40
Conclusion .....	41
Acknowledgements .....	44



“Government policies and programs should view the academic enterprise as a whole and cultivate it as a national asset. This in turn would establish a framework for future partnerships that would allow teaching to flourish in conjunction with research and other activities.”

Mary Lowe Good and Neal Lane, “Producing the Finest Scientists and Engineers for the 21st Century,” *Science* 266 (November 1994).

*The NSFNET Backbone Network Project is managed and coordinated by Merit Network, Inc. under sponsorship of the National Science Foundation, Agreement No. NCR 8720904.*

*This material is based upon work supported by the National Science Foundation under Grant No. NCR 8720904. The Government has certain rights to this material. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.*

## Introduction

Infrastructures, for purposes such as transportation and communication, have long been vital to national welfare. They knit together a country's economy by facilitating the movement of people, products, services, and ideas, and play important roles in national security.

In addition to these broad social and economic uses of infrastructure, however, are individual uses: if directed to do so, the benefits of well-designed and expertly built infrastructure can be realized at the level of the individual or groups of individuals as well as at the level of society as a whole. For over one hundred years, the public switched telephone network has served our needs for communications infrastructure in the United States, reaching most of the population with a



*Figure 1 - Inbound Traffic in Billions of Bytes ,  
Sept. 1991*

ubiquitous, reliable, and easy-to-use telecommunications technology. But today's networking communications technologies can empower scientists, researchers, educators, business people, policy makers, and citizens with the ability to access, create and distribute information much more powerfully and quickly than ever before. Such empowerment represents the true potential of the Information Age.

At present, policy makers, business people, educators, citizen's groups, and individuals in the United States are consumed by the notion of constructing a National Information Infrastructure, or an "information superhighway." Meanwhile, millions of people all over the world are connected to the Internet, a global, packet-switched network of networks running over phone lines and other communications media. As debate continues over the proper means to build, maintain, and oversee such infrastructure, the Internet continues to grow and evolve at astonishing rates.

Thus the development of today's computer networks underlying tomorrow's communications media becomes ever more important technologically, economically, and socially. But in order to understand where we are going, it is vitally important to understand where we have been. The history of the Internet is a tale yet to be told, but the NSFNET represents one of the most significant parts of that story.

Part of the National Science Foundation's ongoing high-speed computing and communications infrastructure initiatives, the NSFNET program from its inception in 1985-1986 was the foundation of the U.S. Internet and the main catalyst for the explosion in computer networking around the world that followed. The NSFNET backbone service, the basis of the larger NSFNET program, linked scientists and educators located on university campuses in the United States to each other and to their counterparts located in universities, laboratories, and research centers all over the world.

Through a cooperative agreement with the National Science Foundation, Merit Network, Inc., a consortium of Michigan universities, in conjunction with IBM, MCI, and the State of Michigan, created the NSFNET backbone service to serve the research and education community's growing

need for networked communications. Using the NSFNET, the research and education community led the way in pioneering new methods of applying information technology to scientific and industrial research as well as education.

The partnership of academia, industry, and government that built the NSFNET backbone service also pioneered a model of technology transfer—the process by which technology is “transferred” from the public to the private sector—whose rewards are only beginning to be realized. From 217 networks connected in July of 1988 to more than 50,000 in April of 1995 when the NSFNET backbone service was retired, the NSFNET’s exponential growth stimulated the expansion of the worldwide Internet and provided a dynamic environment for the development of new communications technologies. As we begin to address the challenge of national and global information infrastructure—the next generation of communications infrastructure—we are fortunate to be guided by the example set by the NSFNET in its successful partnership for high-speed networking.

## I. Coming Together

In 1987, scientists and researchers in some universities and research centers already benefited from access to the first NSFNET, a 56 Kbps backbone network linking five NSF-sponsored supercomputer centers and the National Center for Atmospheric Research, or NCAR. But by mid-1987, traffic on that backbone had increased to the point of overload. “People started to use it, and to like it, and to rely upon it. Their research was already critically dependent on the network being available,” recalls Hans-Werner Braun, co-Principal Investigator of the NSFNET project at Merit, Manager of Internet Engineering at Merit from 1987-1991, and a principal designer/engineer of the routers used on the initial 56K NSFNET backbone. In response to this need for adequate network service, the National Science Foundation issued a *Project Solicitation for Management and Operation of the NSFNET Backbone Network* (NSF 87-37) on June 15, 1987. But who would be served by the creation of an expanded national backbone network?



Figure 2 - Hans-Werner Braun

And who could put such a project together?

The NSFNET partnership team—Merit Network, Inc., IBM, MCI, the State of Michigan, and the National Science Foundation—all brought unique strengths and resources to the project. Each played an important role in building the national backbone network service.

### Merit and the Community to Serve: Research & Education

It may be difficult to believe today, but until recently—within the last two decades—most scientists and researchers were able to carry on their work without recourse to communicating via the marriage of computer and telephone lines known as computer networking. What was it like to be part of the research and education community in the United States before the NSFNET? Scientists had to travel to reach computing resources, from campus to campus or, in many cases, out of the country. Most people didn’t immediately share the results of their research with colleagues outside of their own departments; they would make a telephone call, mail a package, or arrange to meet in person. Researchers and scientists had little or no experience with sending a quick e-mail to notify a colleague of an interesting article, or downloading a data file from a fellow professor in another state, much less working simultaneously on a complex scientific problem in real-time.

Campus networks were increasingly common at major research universities, linking faculty with academic computing centers, but these closed networks were often like islands to one another. This began to change in the United States in the late sixties and early seventies when various wide-area computer networks, or WANs, began operation.

These WANs were primarily Federal research projects, the first of which was the ARPANET in 1969. An outgrowth of the Department of Defense's Advanced Research Projects Agency, the ARPANET's packet-switching scheme was meant to provide reliable communications in the face of nuclear attack. By the early 1980s, other examples of Federal networks included HEPnet and MFENET; however, the variety of different protocols used could impede efficient communication. A small but expanding part of the research and education community used e-mail via UUCP or had access to some kind of message-forwarding system such as BITNET, but access to Federal networking facilities such as the ARPANET was usually limited to computer science departments or restricted to researchers affiliated with laboratories funded by various Federal agencies such as the Department of Defense, the Department of Energy, and NASA.

While for these users, computer networking had become vital to their daily activities, the average professor of English or economics was unaware of the opportunities provided by information technology applications. Ellen Hoffman, manager of the Network Information Services Group at Merit from 1990 to 1993, reminds us:

“We forget that only a short time ago, the Internet was not a word that rolled off anyone's tongue readily. It wasn't so long ago that only the large research universities in the country were connected, and within those universities often there was only one department using the network. Even if there were 50 or 100 schools connected, most people on those campuses had no idea that there was an Internet connection to their campus.”

In the United States, “only about 50 percent of the large research universities had ARPANET connections,” estimates Mike Roberts, vice-president of EDUCOM, the nonprofit consortium formed to promote information technology applications in education.



*Figure 3 - Eric Aupperle, head of Merit Network since 1974, and Sr. Manager in the Univ. of Mich. Information Technology Division*

Merit Network, Inc., the consortium of Michigan universities formed in November of 1966, provided both campus and statewide networking facilities. As a result, it knew, perhaps more than almost any other educational institution, that the lack of a reliable national backbone network was becoming a problem. By the late 1980s, approximately two-thirds of the state universities in Michigan were connected to MichNet, the statewide research and education network operated by Merit. Running its own protocols for

much of that time, the connectivity to the outside world and to each other that MichNet provided instilled a broader understanding among Michigan colleges of the possibilities inherent in “internetworking” compared to most other states. Having seen the value of providing networked communications capabilities to members of the research and education community, Eric Aupperle, president of Merit Network, saw an opportunity to take the fruits of Merit's experience with MichNet beyond Michigan, in an effort that would benefit not only the research and education community but the nation as a whole.

“In an environment of resource constraints, the higher education community has increased its investment in information technology to enhance the

effectiveness of research and education. More and more of the important intellectual output of the higher education community first exists in digitally encoded form. The higher education community will benefit enormously from NSFNET and has a major interest in its development and operation ... NSFNET should be managed by an organization with strong academic roots so that it will be able to anticipate the needs of the academic and research community and build an organization that fits the environment in which NSFNET will operate.” (Merit Proposal, submitted to NSF 14 August 1987)

As part of the research and education community, Merit could provide the foundation for a successful partnership between academia, industry, and government to construct a national backbone network.



*Figure 4: Ellen Hoffman, a member of the original Merit team, managed Merit's Network Information Services group from 1991 to 1993.*

Merit provided the glue that held the project together. That glue was composed of clear objectives, firm yet flexible management, superior staff and the maintenance of close ties to the community served by the NSFNET. The larger currents prompting the initiation of the NSFNET backbone project included evolving Federal networking programs such as the NSF-sponsored supercomputer centers and eventually the High Performance Computing and Communications (HPCC) Program; innovations in computers and networking technology; and the desire to extend networking throughout the research and education community. However, according to Merit, the immediate, overriding objective was simple: “generalized connectivity.” Expanding the 1986 NSFNET, scaling it to support the tremendous growth and interest by the community while recognizing the community's increasing reliance on the network for their day-to-day tasks was Merit's goal.

Another important function Merit provided was to mediate between industry, government and the research and education community.

While the team partners exhibited an outstanding willingness to work together, nonetheless “you're talking about very different cultures,” notes Ellen Hoffman. By having a representative of academia be responsible for the overall management of the project, the partnership was able to avoid competition between industry partners, as well as inspire greater confidence in the NSF and the research and education community that the needs of the community would be well served: “Merit was an excellent organization to have in that situation, because it had the interests of the community at heart,” says Jane Caviness, Program Director for the NSFNET 1987-1990. According to Priscilla Huston, Program Director for the NSFNET 1993-1995, “With the NSFNET program, NSF funded and encouraged Merit, but Merit facilitated the meetings, the open discussion, and got industry and others to be proactive.” But Merit couldn't do it alone.

### **The Role of Industry: IBM and MCI**

Composed of representatives from Michigan universities, the Board of Directors of Merit in 1987 was chaired by Doug Van Houweling, Vice-Provost of Information Technology at the University of Michigan. While working at Carnegie-Mellon University prior to arriving at the University of Michigan, Van Houweling had discussed with IBM an Ethernet-over-satellite communications system for higher education, and suggested that IBM and Carnegie-Mellon make a proposal to the National Science Foundation. Although that didn't happen, as a result of this initial foray, Van Houweling “had been thinking about these issues” when the NSF released its solicitation in 1987:

“In my new role as Chairman of the Board of Directors at Merit, I went back to some of the same people at IBM and suggested this project to them. It took several meetings to get the right group of people in the room, but when we did, there was a lot of common interest.”

Van Houweling got in touch with Eric Aupperle and began a series of meetings with potential industry partners to respond to the solicitation.

Participants in some of these early meetings included Al Weis and Barry Appelman from IBM Research; Bob Mazza, Walter Wiebe, and Rick Boivie from IBM's Academic Information Systems division; and Eric Aupperle and Hans-Werner Braun of Merit. Van Houweling explains the thinking in this way: “We felt we needed a university organization, and Merit certainly filled that role. We needed a software and equipment provider, and IBM certainly fulfilled that role. But we also needed a carrier for the communication,” so Al Weis at IBM talked to Dick Liebhaber, a vice-president at MCI, about the possibility of MCI becoming a partner in the proposal.



*Figure 5: Walter Wiebe (left), IBM Academic Information Systems networking senior manager responsible for the NSFNET program, and Bob Mazza, who led the NSFNET project for IBM.*

According to those involved in these initial talks, there were many reasons why IBM and MCI chose to become part of the Merit NSFNET partnership, including the desire to be part of a high-profile, high technology initiative benefiting the broader university community. However, from a business perspective, becoming involved with the NSFNET backbone service might provide opportunities for “technology transfer.” Broadly defined, technology transfer refers to the process by which the products of government-funded research and development move out of restricted use or research-only application to the private sector.



*Figure 6: Larry Bouman, MCI's Senior Vice-President for Network Operations.*

Executives involved with academic computing within IBM, a world leader in research and commercial products for systems and connectivity, were interested in increased experience and involvement with the Internet protocol TCP/IP, which by 1986 ran on a majority of campus networks. By developing and implementing routers, adapters, network management tools, and other equipment for the NSFNET backbone, IBM would be able learn a great deal, as Al Weis explains:

“IBM was unable to interconnect its large mainframes and some of its new workstations to all the research communities' networks and get adequate performance, because those networks were TCP/IP networks. By working on the NSFNET backbone service, we learned a lot



about TCP/IP and were able to address these needs common in many academic environments.”

MCI similarly grasped that data communications constituted an ever-growing part of the economy, and hoped to develop a reputation for data networking that would match its increasing clout in the long-distance telephone business. The divestiture of AT&T in 1984 in combination with continuing deregulation of the communications industry had created new markets for all kinds of information processing and telecommunications services. An increasing share of corporate, university and government agency budgets was devoted to telecommunications and data processing as the shift from a manufacturing to a service economy took hold in the mid-1980s. With over 34,000 global network miles and a coast-to-coast fiber optic network in 1987, MCI had the ability to provide high-speed circuitry for the NSFNET backbone service and, at the same time, learn about data networking at T1 and greater speeds. Al Weis recalls Dick Liebhaber, a senior vice-president at MCI involved with the NSFNET partnership from the very beginning, saying that MCI had always been known as a voice company, and if MCI was going to learn about data, there was no better opportunity than this project. In this way, two major U.S. corporations signed on to be part of the NSFNET backbone service.

When successful, this “transfer” of technology and know-how is an interactive process, with benefits flowing both ways: industry and academia alike benefit from the increased knowledge and (in the case of the NSFNET backbone) the resulting technology. On the NSFNET project, technology transfer was successful at many levels. Partnering with industry, according to Dale Johnson, the manager of Merit's Network Operations Center (NOC), enabled Merit to “take technology that existed in the labs, and had been proven in the labs, and take procedures that had existed in the business sides of our partners, and bring them into the Internet and make them part of Internet culture.” The NSFNET project helped to distribute technical knowledge and skills among the Internet community, by training personnel in Internet engineering, routing, and network management. Finally, IBM and MCI also were able to gain valuable experience with high-speed networking technology, which later helped them to create commercial products and services based on that experience for the emerging markets for Internet services and products stimulated by the entire NSFNET program. According to Larry Bouman, MCI's Senior Vice-President for Network Operations, “I don't think we'd be where we are today if it weren't for the NSFNET project. Providing the circuitry for the backbone service—and working with IBM and Merit to get it up and running—gave us experience that led to a productive new focus for the company.”



*Figure 7: Harvey Fraser, Project Manager at IBM on-site at Merit for the NSFNET backbone service.*

In return, the NSFNET backbone service's corporate partners invested far more time, money and resources than was required by their joint study agreements with Merit. According to the proposal, IBM would provide the computing systems for the network, including hardware and software, and MCI would provide the underlying telecommunications circuits for the NSFNET. However, by all accounts, there was never a sense of meters ticking away for the partners' contributions, and in fact IBM and MCI were generous with staff time and resources such as financial support or equipment. Harvey Fraser, Project Manager at IBM on-site at Merit for the NSFNET, puts it this way:

“The team worked well because we had resources, executive time, and the desire to make it a success, as well as a lot of good faith and good will. When issues came up, we evaluated them and, if we could, provided whatever was needed. For example, when we needed more routers, IBM got them for us. If we needed more lines for a test network, MCI got them for us.”

Mark Knopper, who succeeded Hans-Werner Braun as Manager of the Internet Engineering Group at Merit in 1991, also remembers that the behavior of NSFNET's industry partners went above and beyond the call of duty:

“I'd go to meetings where the executives would say, 'Well, what are the problems this week?' and 'What can we do to help?' I would be sitting with vice-presidents and top managers of IBM and MCI and they would say, 'What do you need? What can we do to fix things?' Once they knew about a problem, they'd try to find a way to solve it.”

IBM and MCI realized that in working on a project like the NSFNET, which would necessitate the use of uncertain and unproved technology as well as the development of new operating and management systems, a certain amount of risk was involved. The answer was not to simply throw money at the project, but to make a number of different kinds of resources—people, technology, funding—available as needed, according to the judgment of each partner. This relationship of “good faith and good will” among the partners was built upon the high caliber of technical expertise going into the NSFNET project. IBM, MCI and Merit assembled a top-notch team for the NSFNET backbone service: “The quality of the technical contribution from people in Merit, IBM, MCI, all three, was really spectacular,” Doug Van Houweling remembers.

IBM and MCI also brought certain established ways of doing business, “standard operating procedures” forged in the commercial realm that created a strongly production-oriented mind-set for the team. Because of this, according to Dale Johnson, the NSFNET project team “brought applied commercial operating procedures to the Internet. All of this was standard in the telephone companies and also at IBM.”

Documentation, trouble tickets, escalation procedures, “follow-through,” and statistical reports were just some of the many procedures that were adopted by the NSFNET project and modified to be applicable to the Internet. Hans-Werner Braun explains what a difference these procedures made to the thoroughness of project team behavior:

“During one weekend all the routing daemons, all the code that handles routing in the nodes, crashed. I called the person who developed the routing software, Yakov Rekhter at IBM-Yorktown, at home on a Sunday and had a fix for the problem in about two hours. I think

“IBM and MCI essentially donated a lot of resources and a lot of energy for the NSFNET backbone service for the good of the country. Working on a team like that, where people from IBM were typically interested in developing products and moving them into the marketplace, and MCI was typically interested only in selling circuits, it was personally satisfying to see people with different perspectives on what is important in their organizations and their businesses, joining their 'vision' of a national network, so that it could happen. If any one of the partners had balked, we probably couldn't have done what we did.”

—*Elise Gerich, Associate Director for National Networking, Manager of Internet Engineering, Merit Network.*

that's pretty amazing. Also, with MCI, if there was any problem with the lines, they had a very aggressive escalation procedure. Within only a few hours if the problem wasn't resolved it was supposed to be on the desk of a vice-president. It made a huge difference.”

These innovations, which were standard in industry but new to academic networking, would not have been worth much without high-level corporate backing behind them, and according to everyone involved with the NSFNET backbone service, IBM and MCI did indeed provide that commitment. All of the partners—NSF, Merit, IBM, and MCI—demonstrated an overwhelming desire to do whatever was needed to make the NSFNET a success. However, the involvement and dedication of technical and managerial personnel, at every level, at IBM and MCI was among the most important factors in the project's success. With IBM, MCI and Merit, then, industry and academia were well represented for the NSFNET project. The final essential ingredient came from the United States government.

### **The Role of Government: The National Science Foundation**

How did the National Science Foundation become the Federal agency tasked with managing the nation's high-speed backbone network? Other Federal agencies, notably the Department of Defense, the Department of Energy, and NASA, had their own networking projects. According to an influential early article in *Science*, “Computer Networking for Scientists,”(1) in 1986 there were various Federal wide-area networks operating, using protocols from DECnet to SNA and linking hundreds of Federal researchers at sites around the country. Competition between agencies regarding these projects was not uncommon. However, the National Science Foundation had several things in its favor, including the presence of the 1986 version of the NSFNET and support of another network service for computer scientists called CSNET.

The first NSFNET backbone service went online late in 1986, linking six sites, or nodes, together at the five NSF supercomputer centers (San Diego Supercomputer Center, National Center for Supercomputing Applications, Cornell Theory Center, Pittsburgh Supercomputing Center, and the John von Neumann Supercomputer Center) plus the National Center for Atmospheric Research, or NCAR. When the solicitation for the new and improved version of the NSFNET backbone was released in 1987, it sought to build on the success of the 1986 backbone by keeping the basic form of a centrally managed backbone network while expanding the number of backbone sites from six to 13; accelerating the speed (from 56 Kbps to T1 and eventually T3); and improving the reliability and robustness of the service.

NSF staff in the Computer and Information Science and Engineering (CISE) directorate, created on October 1, 1986, envisioned the overall design of the new NSFNET as consisting of a logical topology, or architecture, composed of three “tiers:” the national backbone, various regional networks (such as NYSERNet and Westnet) and campus networks. The overall “internetwork” formation, linking networks of different sizes and running different protocols via TCP/IP, had been borrowed from the ARPANET for the original 1986 NSFNET and would become even more complex with the proposed three-tiered 1987 network.

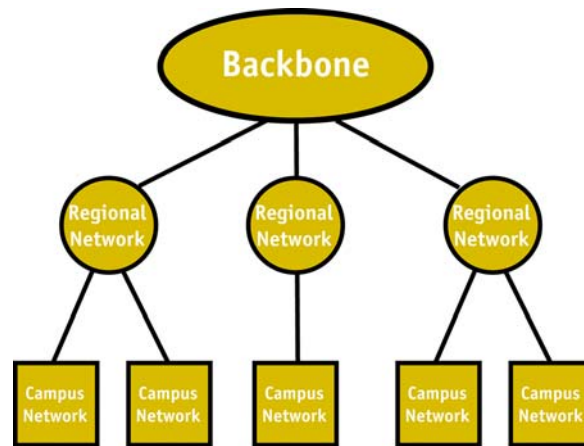


Figure 8: The NSFNET's three-tiered network architecture.

Various organizational and managerial characteristics also distinguished the NSF from other agencies. The NSF's primary duty was perceived to be broader than other agencies. Jane Caviness explains:

“While NSF serves science and engineering generally, other agencies focus on mission-specific needs. The CISE directorate in particular is charged with promoting technology development for all of US research and education.”

Because of this, the NSF also placed fewer restrictions on the use of the NSFNET backbone by researchers and scientists of every stripe, within the limits of NSF's Acceptable Use Policy (AUP) requiring the network to be used in support of research and education. “The way we ran the program made it very easy for people to use the backbone, to connect, to do all kinds of tasks,” says Caviness. Other organizational processes (such as peer review) and instruments (such as the cooperative agreement), while not unique to the NSF, were used to good effect for the NSFNET backbone service. The use of peer review, both in the initial selection and as the project developed, kept the best resources and advice in the networking community available to NSF for the project; it allowed the NSF to stay close to the “cutting edge” of thinking in academia and industry. According to Steve Wolff of NSF, the value of this institutional structure to successful R&D initiatives is undeniable:

“The NSF has been at its best, I think, when it has followed a very rigorous and formal and open policy of peer review. It not only funds the academic community, but it tries to achieve consensus on the worth of what it supports, from the very community that it's supporting. In that sense the community is policing itself. And when the community is a good community, with high standards, it produces some of the very best work that has ever been done.”

Even before the NSFNET, participation in the development of supercomputing and networking initiatives by the academic community had been strong: numerous advisory committees, including the Federal Coordinating Council on Science, Engineering and Technology, and the National Research Council's Computer Science and Technology Board as well as various NSF-sponsored advisory committees contributed ideas, advice and comment to the NSF.

In addition, since the NSFNET was created as a resource for the U.S. research and education community, NSF was able to use a cooperative agreement as the award vehicle. Utilizing the cooperative agreement enabled the NSF to build and maintain a backbone network service in

support of the research and education community while maintaining “a substantial involvement in the process,” says Don Mitchell, Staff Associate at the NSF. A cooperative agreement was chosen as the award instrument for the NSFNET backbone because a contract is used only when procuring specific goods and services for the government, and a grant does not allow the government to significantly guide the management of a project.

The organization within the NSF specifically tasked with overseeing the NSFNET program was the Division of Networking and Communications Research and Infrastructure (NCRI). Foresighted and flexible, NCRI staff seemed to exhibit that most precious commodity of managers: knowing when to step in and when to keep out. NSF participated in the biweekly Partner's Meetings, in addition to the more formal quarterly Executive Committee meetings composed of representatives from each of the partners. But this level of involvement was balanced by a reluctance to micromanage, in addition to the flexibility of the cooperative agreement structure. Hans-Werner Braun puts it this way:

“Having a cooperative agreement with the National Science Foundation, as opposed to a grant or contract, was a real advantage. Often when I needed something done I could call Steve Wolff and say 'Steve, I want to do this, do you have any problem with that?' And he would say, 'No, go ahead.' And I just did it, without any paperwork or additional funding.”



*Figure 9: NSFNET Program Directors (left to right) Priscilla Huston, George Strawn, Doug Gale, Jane Caviness, Steve Wolff.*

The first NSFNET Program Director, and the driving force behind the inception of the NSFNET backbone project, was Dennis Jennings, who came to NSF in 1985 from University College, Dublin. Jennings' immediate successors, Steve Wolff and Jane Caviness, were both closely involved with the NSFNET project. Wolff had been a communications engineer working for the Army research laboratory on various projects in communications and computer systems when he was brought to the National Science Foundation as the second NSFNET Program Director in June of 1986. In

September of 1987, just before the award to Merit was announced, Wolff became Division Director of NCRI and Jane Caviness, Director of Academic Computing Services at the University of Delaware, replaced Wolff as NSFNET Program Director. Caviness held that position until 1990, when she became Deputy Division Director of NCRI. Wolff and Caviness together worked hard to support the NSFNET partners within the NSF, Congress and through two Administrations: “Steve and Jane really have been the driving force,” says Ellen Hoffman. Mark Knopper gives Wolff credit for advancing the cause of the NSFNET:

“The National Science Foundation was very visible and very active in this project. Steve persevered at NSF and at other levels of the government at the time, including then-Senator Al Gore. It was Steve who told them all that we've got to do this. And they said, 'OK'. Steve was able to leverage all of these people and get corporate involvement as well.”

Al Weis puts it this way: “There was one person with a lot of guts, a lot of brains, and a lot of vision. And that person was Steve Wolff.” Additionally, according to Wolff, “Jane really held the program together for a couple of years:”

“When we were building the regional networks, we were empowering them, but with grants that were not enough money to build the networks. Jane was well known in the academic computing community; she was an officer of EDUCOM, and everybody knew her as honest, straightforward, and a hard worker. She went to each regional network, to each person in charge, and said, 'Doesn't matter. We won't be able to get you the money, so you'll have to support it however you can.' And the academic community came through.”

Caviness was succeeded by Doug Gale as Program Director in 1990. Gale did much of the background work for the information services component of the 1993 NSFNET solicitation, and set up new organizational procedures for the NSF “Connections” program. George Strawn succeeded Gale in 1991. Strawn, who was instrumental in realizing the vision for the new NSF architecture, was Program Director until March 1993, when Priscilla Huston, director of Computing Information Services at Rice University, joined the NSF staff. At that time, work had just begun on evaluating proposals for the transition to the new architecture. Huston, a past Chair of EDUCOM's Board of Trustees, was proud of the way the NSFNET project proved the value of increasing access to the NSFNET to the research and education community at large: “There was a very practical reason for introducing networking, but once it was introduced, it became a much more powerful tool than simply in support of supercomputing. Networks began to form the basis for a larger, even international, communications infrastructure,” she recalls.

NSF provided funds for the creation of the expanded NSFNET backbone service in the amount of \$57.9M for the duration of the agreement, originally awarded for five years but extended for a total of seven and one-half years. Although for many years the NSFNET program was the fastest growing program in the NSF in terms of percentage increase in its budget, in comparison to other Federal R&D programs, the NSFNET backbone service, for what it accomplished, represented a relatively modest government investment: “They didn't spend a tremendous amount of money; if you look at the budget of the United States, it was a drop in the bucket,” says Mark Knopper.

The State of Michigan also played an important role in the NSFNET project. Before the solicitation was announced, Doug Van Houweling had a meeting with Jamie Kenworthy, at the time an employee at the state Department of Commerce working on the Michigan Strategic Fund. The Michigan Strategic Fund was tasked with allocating funds from a pool derived from various tax revenues in a way that would benefit the overall development of the economy of the state of Michigan. According to Van Houweling, Kenworthy specifically mentioned that the Fund was interested in initiatives that “combined support from business, education and government” that would affect the future of the state. When it became clear that IBM and MCI would partner with Merit for the proposal, Van Houweling introduced the partnership to Kenworthy and as a result, “one of the very strong parts of the proposal was a letter from Governor Blanchard, in which the Governor essentially pledged that he would use his best efforts to enable an allocation from the Michigan Strategic Fund.” The Strategic Fund contributed \$5M to the NSFNET project, and, according to Van Houweling, “as it turns out, it was strategic for the state of Michigan” because of the impetus Merit provided to the development of networking and Internet connectivity throughout the state.

## **The Vision: Networking the Research and Education Community**

The history of the NSFNET program shows that the overall design and intended uses of the NSFNET proceeded through several stages, from its original conception of a network to link researchers to the NSF-sponsored supercomputer centers all the way to a broad-based networking communications infrastructure connecting the entire research and education community.

Various developments in computer science and networking technology converged with the evolution of Federal programs in computing and communications (as well as similar initiatives in other countries) in the early 1980s to create the context for the NSFNET program, of which the NSFNET backbone service itself was the centerpiece. These factors combined to create the environment and policies that produced the NSFNET and the Internet, and were an important part of the motivation leading the NSFNET partnership to initiate and respond to a project for a national research backbone network. Although the guiding vision of what such services could provide changed over time, in 1987 the NSFNET backbone service solicitation and cooperative agreement were fashioned in response to the perception of specific needs of the research and higher education community for high-speed networking.

High-speed networking was originally conceptualized as a support service for scientific researchers to access high-performance computing resources at the supercomputing centers and elsewhere. In 1985, the primary needs of the scientific research community for a national high-speed network were characterized in the following terms: getting researchers access to supercomputers and large databases, and facilitating collaboration via electronic communication. In fact, to help meet those needs, NSF in 1985 gave DARPA \$4M to install new ARPANET nodes at 40 colleges and universities to be selected by NSF. Steve Wolff remembers, however, that “DARPA had just turned over management and operation of the ARPANET to the Defense Communications Agency, and the bureaucracy was such that it took until 1990 to get all the nodes in place. By that time the T1 NSFNET backbone service had been in use for two years, and the connections to the 56 Kbps ARPANET were redundant. As DARPA decommissioned the ARPANET during 1990, some of its nodes were actually installed and de-installed in the same week.”

In 1986, the five NSF-sponsored supercomputer centers began operation linked by the 56 Kbps NSFNET backbone. This network clearly filled a need: overwhelming demand for networking services quickly saturated the backbone, and the NSF realized that the network would have to be upgraded to support the increase in users.

At the same time, however, individuals within NSF's NCRI as well as the segments of the research and education community that had benefited from other networking research and services (from ARPANET to BITNET) were beginning to express the desire to expand the vision of high-speed networking. Jane Caviness was at the University of Delaware's Academic Computing Services in the early 1980s, and remembers hearing the community's needs described as a “national problem:”

“People were using BITNET and it was expanding, but there were limits to what you could do with it. I was part of a group trying to articulate academic networking needs. We realized that this was a national problem that had to be addressed, and the community was organizing to express that need. A set of groups came together at that time to help get the effort started.”

“NSF started to create operational TCP/IP infrastructure initially to interconnect the supercomputer centers, but very shortly thereafter the vision of the NSFNET backbone service

was broadened to include the research and education clientele at large,” says Hans-Werner Braun. “In fact, shortly after Steve Wolff started at NSF he made comments to me, and I assume to others as well, saying ‘I do not want to see this network only as a supercomputing center network.’”

A paper released in 1987 by EDUCOM noted that “as limited budgets force us to examine closely the means to improve research and scholarship, information technology emerges as a key tool ... faculty in all disciplines have a need to use networks for communications among their peers and within research projects.”(2) Thus, while NSF staff at NCRI was primarily concerned with expanding the design and operational aspects of the NSFNET backbone at the time of the 1987 solicitation, they were already planning to extend the vision of the NSFNET program by attempting to build a more comprehensive infrastructure and by encouraging all of the higher research and education community to connect.

The vision for national high-speed networking in 1987, then, contained two main elements. Within NCRI, there was a growing consensus that the best way to create a ubiquitous national networking program was to build infrastructure on many levels at once. Over time, the NSFNET program would include funds to construct, operate and maintain the NSFNET backbone network, to support the regional and midlevel networks, and to fund access to the NSFNET for individual colleges and universities. According to Steve Wolff, following this strategy was the best way to build the network in a way to benefit the most people.

“The lucky thing that we did was to say why don't we just make a cross-country transit network, and then empower the regionals? We'll do the backbone, we can manage that, but we'll let the regionals connect everybody else. In this way NSF empowered hundreds of people all at once.”

Accordingly, the NSF's first goal for the new NSFNET was to construct a national backbone network service, linking the supercomputer centers and the emerging regional and midlevel networks at T1 speeds. Secondly, expanding use of the network from scientists and researchers to the general academic community would also coincide with the growing understanding that in America's universities, research and educational functions were thought to benefit one another. New networking technologies might further erase the boundaries between research and teaching, and even between teacher and student. In this environment, then, separating “research” use of a network from “educational” use would not only be technically difficult, but it didn't seem to make much sense, not when networking the entire research and education community could improve the learning process at all levels. According to Priscilla Huston,

“Very early on, both the campus communities and the government realized that it would be very difficult in the networking world to separate one thing from the other. Historically we have tended to think of teaching and education as separate from research. We have always thought there was value in the teacher being a researcher as well as a professor, but networking has facilitated a broader learning experience where students are empowered, students learn better, and may in fact make contributions themselves. In this way, networking can help to foster a more interactive learning process.”

During 1987, the specific vision for the NSFNET backbone solicitation began to take shape, taking into account this expression of the community's needs.

## **The Solicitation**



As the currents of opinion regarding networking for the research and education community swirled around computer scientists, scientific researchers of every stripe, government policy makers, and members of the computer and communications industries, the NSF's specific, immediate needs in creating Solicitation 87-37 centered around expanding and building upon the 1986 NSFNET. As more of the scientific research community became connected, they realized just how helpful networking could be, and this encouraged a dynamic process of innovation in networking. New applications such as NCSA Telnet were written to take advantage of the growing information resources available on what was becoming the Internet. This in turn attracted more users to the NSFNET. Moreover, NSFNET users increasingly incorporated networking into their daily lives, relying on their ability to communicate electronically in all aspects of their work.

Reflecting this increasing dependence on networking, one of the basic rationales behind the creation of the NSFNET program was beginning to shift. Rather than creating and supporting a “research” network for computer scientists and networking researchers to test new networking technologies, as it had been envisioned in the early 80s, by the time of the 1987 solicitation the new NSFNET backbone was conceived to be a “production” network, meaning that the network would be built, operated and maintained at a high level of performance and reliability. According to Hans-Werner Braun, the NSF had high service expectations for the network, and did not intend the NSFNET to be used only for research:

“The NSFNET backbone was really there to provide infrastructure. Some of the networking researchers did not understand even until 1988-89 that we were operating a network with very high service expectations that people really depended on. They needed to be sure that the network was always up and that networking researchers didn't experiment with it.”

The NSF also realized that in order to scale the network backbone to the level demanded by its exceptionally high growth rates, significantly more investment would be needed. Thus the solicitation itself was drafted to encourage participation by companies like DEC, AT&T, Sprint, IBM, and MCI, by not specifically excluding private companies from bidding, and by framing the technical requirements so that by meeting them, companies would gain experience and have a better competitive position in the future.

Campuses spent approximately \$100,000 per year per campus for connectivity to the 1986 incarnation of the ARPANET, according to EDUCOM, but other significant costs included extensive installed bases of incompatible computing systems, which often had to be adapted or replaced. The NSF knew that it would not be able to fund a new, expanded NSFNET in its entirety; hopefully, the careful administration of funds to various levels of the NSFNET (the backbone, the regionals, and campus network connections) would seed the development of the overall program, while cost-sharing arrangements with industry partners would provide the remainder of the necessary investment.

The solicitation was drafted at the NSF by Steve Wolff and Dan Van Bellegham. Funding provided by NASA made it possible for Steve Goldstein, at the time an employee of MITRE Corporation, to assist by preparing an analysis of prospective costs and performance of the new backbone. One of the firm requirements of the solicitation concerned the communications protocol standard: it had to be TCP/IP, created by Vint Cerf and Bob Kahn in 1973 as a research project on behalf of ARPA. Dennis Jennings chose TCP/IP, as opposed to other protocols, for the new NSFNET backbone because it was an open and non-proprietary standard, and also because the ARPANET used TCP/IP: “One of the things that helped jump-start the NSFNET was an

agreement with ARPA to allow us to use the ARPANET for NSFNET traffic,” recalls Jane Caviness.

In addition to the communications protocols, appropriate hardware and software had to be specified, which was difficult at the time because there were few, if any, off-the-shelf TCP/IP products available. The fastest routers available at the time switched fewer than 1,000 packets per second, compared to today's high-speed routers with performance levels of up to 100,000 packets per second. An advisory panel of scientists and engineers from academia and industry, including those involved in CSNET, assisted NCRI staff in preparing the solicitation.

Finally, in an open meeting at the NSF's offices on June 15, 1987, the solicitation was released. Its architects now had only to wait. For those interested in accepting the challenge put forth by the NSF, however, the long journey toward a national research and education backbone network had only begun.

### **Drafting the Proposal and Winning the Award**

For about six weeks before the due date for proposals, a group of people from MCI, IBM, and Merit met at the University of Michigan, Merit's administrative home, to put ideas together: Eric Aupperle, Hans-Werner Braun, Walter Wiebe, Paul Bosco, Rick Boivie, S.S. Soo, Mathew Dovens, Bob Mazza, Len Dedo, Doug Van Houweling, and Greg Marks led the proposal development team.

“It was exciting stuff, hammering the proposal together. There were a lot of bright people and a lot of enthusiasm to go with it; we burned both ends of the candle on that one,” recalls Mathew Dovens. Walter Wiebe and Paul Bosco worked constantly within IBM to solicit the necessary commitment: “For most of the summer the phone was glued to my ear, begging, praising, eliciting, humbling, whatever it took to get whatever we needed from within IBM to move forward,” recalls Paul Bosco, manager of the team at IBM responsible for building routing systems, adapters, and technologies for the NSFNET backbone service.

At the IBM senior executive level, too, persuasion was the order of the day. Bob Mazza, who led the NSFNET project for IBM, found himself on the phone with vacationing executives, selling them on the proposal. “We were in a period of major changes and cost-cutting, so it was a particularly tough time to get an agreement,” he recalls. But IBM had a long history of interest in supporting the academic community, with influential joint studies already in place with MIT, Carnegie-Mellon University, and the Cornell Supercomputing Center. “The NSFNET backbone service was an ideal adjunct to those projects,” says Mazza.

At the same time that the partners-to-be had definite ideas about how to respond to the NSF solicitation, making all of the pieces fit properly into a cohesive, integrated design for a backbone network service required them to break new conceptual ground. There were three main components of the basic architecture in the proposal: the packet-switching nodes, called the Nodal Switching Subsystem, the circuit switches interconnecting these nodes, called the Wide Area Communications Subsystem, and a network management system.

In addition to overall engineering, management, and operation of the project, Merit would be responsible for developing user support and information services. IBM would provide the hardware and software for the packet-switching network and network management, while MCI would provide the transmission circuits for the NSFNET backbone, including reduced tariffs for that service.

A team from IBM, MCI, Merit, and the University of Michigan labored for many late nights to write the proposal, which was completed and submitted to the NSF on August 14, 1987.

As the NSF reviewed the proposals, the team members had no idea if they would win. Through the fall of 1987, however, the partners did not rest while waiting for the award announcement. Anticipating the NSF's desire for more concrete examples of the potential of the Merit proposal, IBM built a simulation network to test some of the hardware and software systems to be used in the NSFNET backbone service. The partners also worked on development of hardware and software systems for the packet switches to be located at each node, as well as testing of the proposed fourteen MCI T1 circuits. On November 24, 1987, the NSF announced that they had a winner: Merit and its industry partners, together with the State of Michigan's additional funds, would build a new NSFNET backbone.

The new Cooperative Agreement NSF 872-0904 was announced in a press conference convened at Wayne State University, one of Merit's member campuses, and a press breakfast in Washington D.C. for members of Congress, both of which were attended by representatives of the new NSFNET project team. At that time, networking still represented barely a blip on the radar screen of most news media; those members of the press that could be cajoled to come received something of a crash course on "internetworking," an introductory lesson to be repeated to ever-expanding groups of people over the course of the award. Today, of course, the words "information superhighway" make front-page news in national newspapers and magazines. Back in 1987, however, Merit and its partners were charged with planting the seeds of that future.

## II. The Partnership

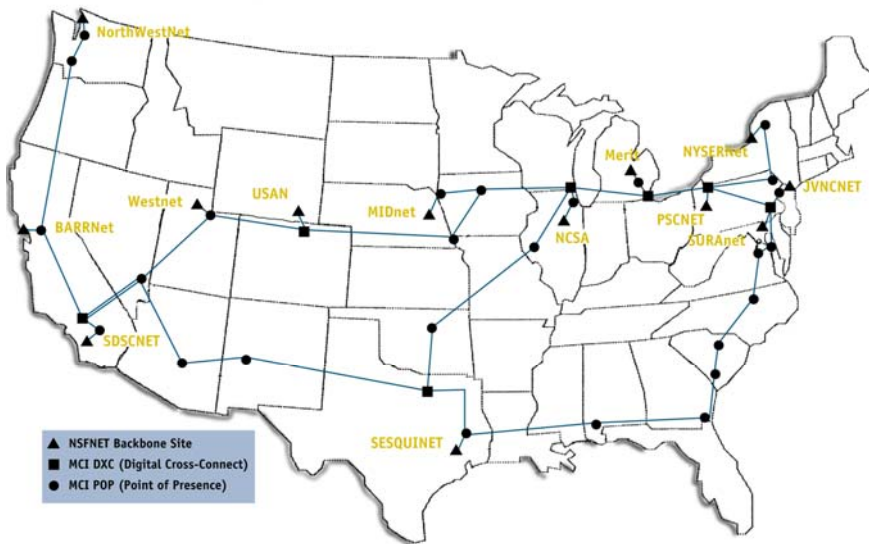


Figure 10: Physical Plan for the NSFNET T1 Data Network

From the moment the award was announced in November of 1987, there was only one goal for the partnership: building a high-speed national network backbone service. No one had ever attempted a data networking project of this scale, with backbone, regionals, and campus networks forming such a large internetworking environment. Merit, under cooperative agreement to the NSF, had responsibility for design, engineering, management, and operation of the NSFNET backbone service, but the project required all of the partners to work together to respond to enormous technical and managerial challenges.

Operational in July of 1988, a scant eight months after the award, the new T1 backbone linked thirteen sites, transmitting 152 million packets per month at a speed of 1.5 Mbps. Merit, IBM and MCI also developed a state-of-the-art Network Operations Center (NOC), and staffed it to run 24 hours a day, seven days a week; it was one of the first of its kind on the Internet. By 1989, Merit was already planning for the upgrade of the NSFNET backbone service to T3 (45 Mbps). In such a dynamic environment, the partnership had to walk a delicate line between foresighted planning and on-the-go engineering and user support. The tremendous success of their efforts is testament to the dedication and talent of all the members of the project team.

### Building a National Backbone Service

Merit had promised to deliver the NSFNET backbone service by July 1, 1988. “No one thought that it would be up in such a short period of time, but we did it,” remembers Jessica Yu, a member of Merit’s Internet Engineering Group. “The Internet engineers, only a handful of us at the time, worked day and night to build the network. The work was intense, but we didn’t see it as a burden. It was very exciting; we felt that we were making history with the project.”

The first thing Merit needed to do internally was to pull together the group of people that could make it happen. In the mid-80s, “there were few people who knew about this kind of technology,” says Hans-Werner Braun. Not only was it difficult to find people with the necessary background; it was made even more difficult due to the tight deadline for start-up of the backbone. Braun, who had been working on TCP/IP for years, says: “The NSFNET program

educated a generation of networking people, directly as well as indirectly, by making all of this technology available.”



*Figure 11: Elise Gerich, Merit Associate Director for National Networking, was a member of the original Merit team, and has been Manager of Merit's Internet Engineering Group since 1994.*

After the Internet Engineering team at Merit was assembled—Braun, Bilal Chinoy, Elise Gerich, Sue Hares, Dave Katz, and Yu—they established contacts with their counterparts at IBM, MCI, and the regionals, relationships that fostered continuous communication among the technical team in much the same way as relationships between executives such as Doug Van Houweling, Eric Aupperle, Dick Liebhaber, Bob Mazza, and Al Weis had provided the contacts for the initial partnership.



*Figure 12: Members of the Merit team (left to right): Dale Johnson, Sue Hares, Eric Aupperle, Ellen Hoffman, Mark Knopper, Jessica Yu.*

In addition, a special engineering team was assembled by IBM and placed at the Merit offices in Ann Arbor. The group was led by Jack Drescher, then a Product Manager at IBM's Research Triangle Park Laboratory, and included Jim Sheridan and Rick Ueberroth. Drescher, working in Michigan, Walter Wiebe and Paul Bosco in Connecticut, and Jordan Becker in New York pulled together and managed the detailed plan of how the partners were going to get from February of 1988 to online production in July of 1988, and how the system would be improved after that. Later, in mid-1989, Harvey Fraser took over as Manager of the on-site team.



*Figure 13: Jack Drescher, who headed IBM's special engineering team at Merit.*

For the initial deployment of the NSFNET backbone service, there was a lot of talk about getting the machines to the sites and getting them configured. We thought we were going to have to send people around to each site,” recalls Elise Gerich, now Manager of Internet Engineering and at the time site liaison at Merit for the regional networks. Instead, Drescher's IBM team and Merit came up with the notion of creating a depot or assembly line, in the manner of Henry Ford and his first automobiles, at Merit in Ann Arbor.

“We turned the third floor of the Computer Center into an assembly line and delivered all of the IBM RTs there. Jim Sheridan from IBM and Bilal Chinoy from Merit spent a couple of weeks testing all the machines and installing all of the software. We even sent out notes to the various regionals saying that if you want to come to Ann Arbor, we would be happy to put you up and you can build your own router. In the end we had the whole floor covered with parts and machines and boxes; it was a great way to deploy everything.”

Staff from the regionals did, in fact, come to Ann Arbor to help in the assembly depot, where hardware and software for the backbone nodes (Nodal Switching Subsystems) were assembled and configured. Each of the thirteen NSSs was composed of nine IBM RT systems processors running the Berkeley UNIX operating system, linked by two token rings with an Ethernet interface to attached networks. The nodes provided packet switching (i.e., forwarding packets received from one link to another link to their destination), routing control (directing packets to follow a “route” along the links to their destination), and statistics gathering (for network traffic information) for the network. Redundancy was an important part of the design: if a component or an entire system failed, another RT would be available to take over for it.

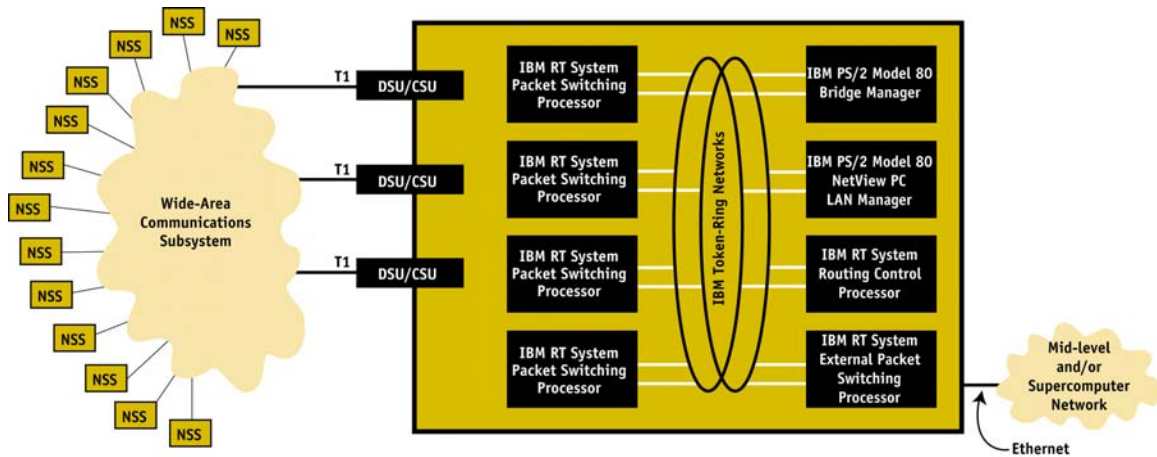


Figure 14: Redundancy was an important part of the Nodal Switching Subsystem design. The three IBM RT's that served as backup are not shown in the diagram.

Each NSS, then, was essentially a large packet switch for transmission of data between regional and midlevel networks and the NSFNET backbone service. The new T1 NSFNET would link 13 sites: Merit, the NSF supercomputer centers plus the National Center for Atmospheric Research, and the following regional and midlevel networks: BARRNet, MIDNET, Westnet, NorthWestNet, SESQUINET, and SURAnet. NYSERNet and JVNcnet were also served by the backbone, because each was collocated at a supercomputer center.

IBM service personnel traveled to each midlevel and regional site to install the new systems, with remote support from MCI and Merit. “It was an amazing team effort,” recalls Walter Wiebe, the IBM Academic Information Systems networking senior manager responsible for the NSFNET program, including hardware, software, network engineering, and field support. “Everybody worked together extremely well—designing, integrating, testing, and documenting the hardware and software, and configuring 150 systems that integrated thousands of parts.”



Even with the Merit “depot,” however, there were occasional snags. “While shipping a node to Rice University for the SESQUINET connection, it turned out that they didn't have a delivery dock. The truck driver got there, took the equipment off the truck, left it in the driveway, and told the SESQUINET staff, 'Well, you're going to have to get it inside,' up the stairs and through a hallway,” Gerich recalls with amusement. “There were seven RTs and three racks, so it weighed a lot.” Making sure the truck driver stayed on hand, SESQUINET contacted Gerich, who made a few telephone calls to ensure that the driver got instructions for inside delivery. “It was so funny, that he was just going to abandon it outside the building,” says Gerich, illustrating just how new and unfamiliar building a national backbone network on that scale was to many parts of the research and education community.

For the new Network Operations Center, the same third floor of the Computing Center at the University of Michigan that had been used to assemble the nodes was again transformed, this time into a showcase for the “nerve center” of the backbone. The center of the NOC was a hexagon-shaped room with large monitors showing the status of each NSS, the statewide network, and the NSFNET links. Through a large window on one hallway one could see the machine room with the Ann Arbor NSS as well as the two IBM 4831 mainframes. These machines handled network management, statistics collection and also acted as a server for NSFNET information such as documentation and data on network growth and performance, usage of the backbone, and distribution of networks.



*Figure15: The round-the-clock NOC set new standards for coverage within the Internet.*

Soon after the NOC was constructed, it was staffed to support the NSFNET backbone service twenty-four hours a day, seven days a week with an increase from four to eighteen people. Besides the NSFNET backbone, the NOC was also responsible for network operations for the University of Michigan's campus network, MichNet, and CICNet, a regional network linking the Big Ten schools. “We were funded to build a commercial NOC long before anybody else, so we were able to provide levels of service that astounded the funds-poor community at the time,” says NOC manager Dale Johnson. As the backbone was developed, operational issues, including troubleshooting functions, moved out of the Internet Engineering group at Merit into the NOC.

Inventing a nationwide backbone network service from scratch, in the context of increasing use and constant change, meant not only that new operating procedures had to be devised, but if something wasn't working, a solution had to be found immediately. NOC operators often called the Internet Engineering staff in the middle of the night to fix problems.

The NOC team developed an automated trouble-ticket system for use among the partners, including the escalation procedures adapted from MCI; all of these procedures were eventually collected into an NSFNET Site Manual for use at other network NOCs. The team also created methods of gathering network traffic statistics and forwarding them into reports which were made publicly available on nic.merit.edu, Merit's host system for information services. In addition, network management, it was realized, was thought of a bit differently by IBM, MCI and Merit; all of the partners had to exchange information and learn from one another's experiences. Elise Gerich gives an example:

“We were working hand-in-hand with Bill Spano, Ken Zoszack, and others from MCI who'd never dealt with data networking; they'd only had experience with voice networks. We'd call them up and say, 'look, we're seeing that you guys are dropping packets,' they would run tests on their telephone lines and say that the traffic met all of their criteria. We'd reply that we were only getting 50% of our packets through, and they would say, 'The line looks good to us.' So it was an education process: they were learning how our tools could show there was something underneath in the physical layer that wasn't working. MCI had never looked at their circuits that way before.”

Quick access to the IBM team in Ann Arbor and the MCI staff at its offices in Southfield, Michigan, was a boon to Merit. Jack Drescher says “many were the times when Hans-Werner and Bilal would hustle down the hall to my office and say 'We think we have this problem ...' or 'Did you know about this problem?' We'd sit down and jointly work out an action plan—it was great teamwork.”

Software tools detailed in the proposal for network management and routing had to be revised and adapted to the new environment. “As it turned out, some of the technologies we expected to work did not,” explains Paul Bosco.

“Instead of being interoperable, various TCP/IP implementations failed. At times Macintoshes couldn't talk to Sun workstations, IBM machines couldn't talk to DEC machines, and nothing could talk to PCs. But it turned out that many of the problems were related to the deployment of TCP/IP stacks on 'real' wide-area networks for the first time.”

To provide a place to explore some of these problems as well as develop new solutions, early on in the first year of the project the team members deployed a test network between Merit in Ann Arbor; IBM in Milford, CT and Yorktown Heights, NY; and MCI at Reston, VA for research, development and testing. “MCI does what many large corporations do: we perform extensive testing on new equipment and software before we put it in a production environment. What NSFNET demanded was a bit more flexibility,” explains Mathew Dovens, recalling how the team members pulled together when problems occurred in the network. Bob Mazza looks back on the partners' decision to build the test network as instrumental to the success of the NSFNET backbone service. “The power of the test net was that it ran in parallel with the backbone. Being able to test changes before deploying them on the production network gave the team a real edge in keeping up with all the new technical requirements.”



**T1 NSFNET Technical and  
Administrative Contacts:  
July 1988**

**BARRNet (Palo Alto CA)**

Abraham Bleiberg, Tom Ferrin, Phil  
Urquiza, David Wasley, Bill Yundt

**JVNCnet (Princeton NJ)**

Joanne Boyle, Sergio Heker,  
Dan Leonard

**Merit (Ann Arbor MI)**

Eric Aupperle, Hans-Werner Braun

**MIDnet (Lincoln NE)**

Dale Finkelson, Doug Gale,  
Gerald Kutish, Mark Meyer

**NCAR/USAN (Boulder CO)**

Britt Bassett, Joe Choy, Don Morris

**NCSA (Champaign IL)**

Sue Greenberg, Ed Krol, Ross Veach

**NorthWestNet (Seattle WA)**

Marian Folkner, Steve Hallstom,  
Ronald Johnson, Dick Markwood

**NYSERNet (Ithaca NY)**

Scott Brim, Craig Callinan, Mark Fedor,

Jeff Honig, Richard Mandelbaum,

Marty Schoffstall, Bill Schrader

**Pittsburgh Supercomputing Center**

Jim Ellis, Gene Hastings, Mike Levine,

Dave O'Leary, Marty Schulman

**San Diego Supercomputer Center**

Jay Dombrowski, Daniel Drobnis,

Paul Love, Gerard Newman

**SEQUINET (Houston TX)**

Guy Almes, Farrell Gerbode

**SURAnet (College Park MD)**

Jack Hahn, Mark Oros

**Westnet (Salt Lake City UT: Boulder CO)**

Pat Burns, Allen Cole, Mike Moravan,

David Wood

Once the hardware, software, and circuits were in place at each of the backbone nodes, Merit began the complicated process of integrating all the components, debugging code, and configuring the links between each site, so that traffic could begin to traverse the new backbone. Merit also began to work with regional engineers to prepare for the cutover to the new T1 NSFNET backbone service. Jessica Yu remembers that “the challenge was to get the regionals to move to the backbone.”

“The network was so new, and we put it together in such a short time, that the midlevel networks were reluctant to move their traffic onto it. We decided that the first network we'd move would be MichNet, then called the Merit Network. When the regionals saw that the new service was solid, they cut over to it one after another.”

Merit and IBM's careful routing design and engineering work, led by Yu and IBM's Yakov Rekhter, along with close collaboration with network staff at the regionals, helped ensure a smooth transition to the T1 NSFNET. The cutovers were scheduled for late nights or early mornings, to keep service disruptions to a minimum. “It was a lot of work,” according to Yu, “but it helped us build a trusting working relationship with the regional network engineers that was essential for network operations, and continues today.”

In addition, Merit began to publicize the impending changeover to the new backbone.

“There were things that needed to be done to support the sites that were coming up, and to make people aware of it—not just create a technical showcase, but make sure it was showcased,” says Ellen Hoffman. Visitors to the new NOC received full tours; Merit also produced slide shows and information packets for the many presentations to be made about the NSFNET over the coming months and years. Merit's publishing activities ranged from engineering and site liaison group working papers and technical documents, to a biweekly newsletter called the *Link Letter*, which went to the regional and midlevel networks and other interested readers.

In July of 1988, the new NSFNET backbone service, with over 170 networks connected, went online. For the first time, the research and education community could access a high-speed, high-

performance communications and data network backbone service, and they responded in the most direct way possible: within a month, network traffic doubled on the NSFNET backbone. During the first month of operation, traffic from the 56 Kbps NSFNET was phased in gradually in order to monitor the transition. In a letter to Eric Aupperle from Guy Almes, at the time chairman of the Federation of American Research and Education Networks (FARNET), the NSFNET team was praised for its achievement:

“We have seldom seen a major networking project come off so smoothly, and never one of such a magnitude. The hardware was new, the software configuration was complex and innovative in a number of ways, the line speeds were an order of magnitude faster than prior NSFNET speeds, and the loads borne during the first month of operation were heavy. Despite these factors, the NSFNET backbone came up on schedule with high performance and reliability even during the first weeks of use.”

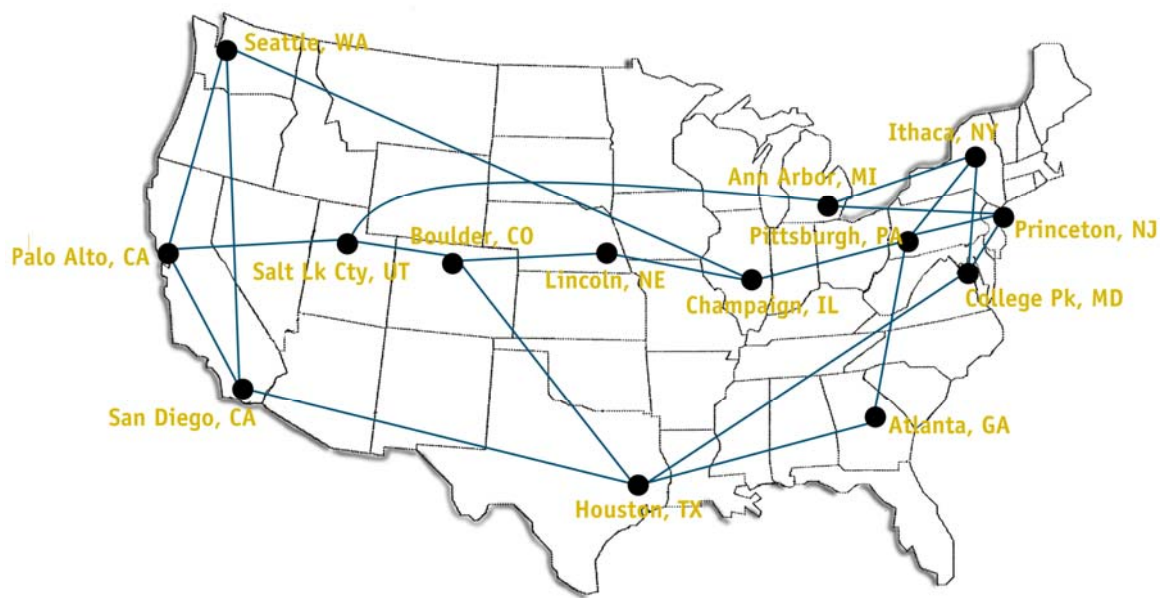


Figure 8: The Expanded T1 Network. An additional Atlanta node, deployed in 1990, brought to 14 the number of regional networks, midlevel networks, and supercomputer centers connected to the NSFNET at T1 speeds.

The NSFNET partnership, which had put in unbelievably long hours in order to meet the July start-up date, had performed beyond anyone's expectations, including perhaps even their own. By July 24, 1988 the old NSFNET backbone service, with its routers showing signs of endless code-tweaking and weary circuitry, was honorably decommissioned. The new NSFNET backbone service had arrived.

## Transition to T-3: Challenges and Opportunities

During the first three years of the project, NSFNET team members focused on building and improving the NSFNET T1 backbone network service. In 1989 the backbone was reengineered, increasing the number of T1 circuits so that each site had redundant connections to the NSFNET backbone as well as increasing router capability to full T1 switching. These changes significantly increased the robustness and reliability of the backbone service. IBM and Merit also continued to develop new network management and routing tools.

### ***Initial Contacts at the NSFNET Expansion Sites***

#### **NEARNET (Cambridge MA)**

John Curran, Dan Long, John Rugo

#### **Argonne National Lab (Argonne IL)**

Larry Amiot, Linda Winkler

#### **DARPA (Arlington VA)**

Paul Mockapetris, Mike St. Johns

#### **FIX-East (College Park MD)**

Jack Hahn, Mark Oros

#### **FIX-West (Moffett Field CA)**

Milo Medin

#### **SURAnet (Georgia Institute of Technology Atlanta GA)**

Al Hoover, Dave O'Leary

Perhaps the most concrete example of the NSFNET's success, however, was the exploding growth in usage by the research and education community. Even though the applications, like telnet, e-mail and FTP, were not necessarily the most "showy," once connected, scientists and researchers didn't know how they had ever gotten along without it. Their enthusiasm led to over 10% monthly increases in usage of the NSFNET backbone, a rate the overall Internet has continued to experience to this day.

"When we first started producing those traffic charts, they all showed the same thing: up and up and up! You probably could see a hundred of these, and the chart is always the same," recalls

Ellen Hoffman, "whether it is growth on the Web, growth of traffic on the Internet, or growth of traffic on the regionals. You didn't think it would keep doing that forever, and it did. It just never stopped."

By 1989, the tremendous growth in use of the NSFNET prompted the NSF and the other team members to think about expanding the backbone. The traffic load on the backbone had increased to just over 500 million packets per month, representing a 500% increase in only one year. Every seven months, traffic on the backbone doubled, and this exponential growth rate created enormous challenges for the NSFNET team.

The NSF also had encouraged greater use of the network by funding many more universities and colleges to connect to the NSFNET through its "Connections" program. IBM and Merit conducted tests based on usage statistics, and projected that the NSSs on the T1 network would saturate by 1990. An 18-month performance review by an outside committee had pronounced the team's performance excellent, and this encouraged the partners to develop a working plan for the upgrade. In view of this continuing demand, the NSFNET partners felt that exercising the option to deploy a higher-capacity network was the right thing to do.

In January of 1989, in an Executive Committee meeting, the partners introduced their plan to upgrade the backbone network service to T3. The NSF also wished to add a number of new backbone nodes, so Merit was requested to prepare proposals for the added cost of new backbone node sites at T1 and T3 speeds, while the NSF issued a solicitation to the community for those interested in becoming new NSFNET node sites. It was eventually decided by the NSF that the team would implement eight new T3 nodes and one new T1 site. Of the eight T3 nodes, six of

the original T1 nodes would be upgraded, and two new T3 node sites would be added. Finally, the remaining T1 nodes would be upgraded to T3, increasing the total number of backbone nodes on the NSFNET from 13 to 16, all running at 45 Mbps.

The first step was to upgrade the test network to T3, using collocated T3-capable NSSs and T3 circuits. Throughout 1989 and 1990, the partners, particularly IBM, strove to implement the brand-new technology on the test network. Central to the new network were the improved node systems.

“We built the first robust, high-performance T3 routers operational anywhere in the world,” says Alan Baratz, Director of IBM’s High-Performance Computing and Communications activities within IBM Research. Each T3 node consisted of an IBM RS/6000 workstation that performed many of the same functions as the nine IBM RTs that made up each NSS—but in a single box, and at a much higher speed.



*Figure 9: Alan Baratz, Director of IBM's High-Performance Computing and Communications activities within IBM Research.*

Two types of T3 routers were deployed on the backbone. Core Nodal Switching Subsystems (CNSSs), located at junction points on MCI’s network, enabled a much closer match between MCI’s circuit-switched network and the NSFNET backbone’s packet-switched network. Exterior Nodal Switching Subsystems (ENSSs) were installed at regional networks attached to the NSFNET, and acted as end nodes for the backbone. These next-generation packet switches in the nodes would eventually switch over 100,000 packets per second.

“In essence, it was a matter of developing some new, leading-edge technology that utilized some of the principles of parallel processing, where each adapter card, running a subset of the UNIX operating system and its own copy of the IP protocol, forwarded packets at very high speeds,” according to Rick Boivie, NSFNET Development Manager at IBM.



*Figure 10: Rick Boivie, NSFNET Development Manager at IBM.*

Another important part of the new T3 backbone service was MCI’s digital cross-connect systems, which made the routing, management and (in case of trouble) restoring of dedicated circuits “as simple as sending commands to electronic switches in the MCI junctions,” explains Mathew Dovens.

At Net '90 in Washington D.C., a yearly networking conference sponsored primarily by EDUCOM, the NSFNET partners exhibited a prototype T3 service in a ground-breaking demonstration, linking the conference site to Merit and the NSFNET backbone at 45 Mbps. At the press conference, Steve Wolff, Doug Van Houweling, Mike Connors, at the time Director of Computing Systems at IBM, Dick Liebhaber, and Jamie Kenworthy from the State of Michigan

spoke to national print and broadcast media representatives about the NSFNET project and the proposed T3 upgrade.

Upgrading the NSFNET backbone service to T3 provided both challenges and opportunities to the project team members as well as the research and education community they served. The tremendous advances required to develop and implement the next generation of high-speed networking would push the envelope of technology yet again, and if successful would raise the level of performance of the network another giant notch, enabling more and greater capabilities for the users of the network. But the challenges and opportunities were not only technical in nature. It was becoming clear to the partners that to the extent that the community relied upon a robust, high-performance backbone of greater and greater capacity, more investment would be needed to satisfy the demand. In addition, non-academic organizations willing to pay commercial prices increasingly desired Internet connectivity, but were restricted from using the NSFNET backbone due to the NSF's Acceptable Use Policy, which defined the "research and education traffic that may properly be conveyed over [the] NSFNET [backbone service]."(3)

This growing use of the Internet for purposes other than research and education, together with the challenges presented by an upgrade of the backbone service to T3, precipitated an opportunity of another kind: the next stage of technology transfer and the beginning of the Internet industry.

In September of 1990, the NSFNET team members announced the creation of a new, independent nonprofit corporation, Advanced Network & Services, Inc. Al Weis, president of ANS, explains the thinking behind the formation of ANS:

"No one had ever built a T3 network before; there were no commercial T3 routers available and telephone companies had little experience with clear-channel T3 lines. I knew what IBM had and didn't have, and I realized that if we were going to make this jump to 45 Mbps, with the network constantly growing, it was going to be a very difficult jump. To do this, you had to have an organization that was technically very strong and was run with the vigor of industry."



*Figure 19: At the Washington, D.C. announcement of ANS were (left to right) Doug Van Houweling (Merit), Al Weis (ANS), Dick Liebhaber (MCI), and John Armstrong (IBM.)*

According to Weis, the commitment to commercial provision of high-speed networking would attract corporate customers, which would in turn provide more funds to support the backbone from which the research and education community benefited. From now on, ANS would provide the NSFNET backbone service as a subcontractor to Merit, and would also take over the Network Operations Center. IBM and MCI initially contributed \$4M to ANS, with an additional \$1M pledged from each, as well as personnel and equipment; Merit was represented by Doug Van Houweling on the Board of Directors of ANS. In May of 1991, ANS spun off a for-profit subsidiary called ANS CO+RE Systems, "so that if we did anything that was commercial and

taxable, we could pay tax on it,” says Weis. “Both IBM and MCI felt a need to spin off an organization that could truly focus on evolving the network further,” says Matthew Dovens, who became the MCI on-site liaison to ANS.

Team members, including ANS, began the painstaking process of installing, configuring and debugging the new nodes. Almost from the start, they ran into difficulties. Doug Van Houweling describes the enormity of the technical challenges:

“The T1 network required that we build new packet switches, but the fundamental T1 data transmission technology was pretty solidly in place. What we didn't understand when we went to T3 was that we not only had to do the switches all over again, but we also had to learn how to transmit data over a full T3 line, which wasn't being done. So when we built the T3 network we had two technology frontiers to overcome, and they interacted with one another in all kinds of vicious ways.”

Throughout 1991 several T3 nodes were installed, but production traffic remained on the T1 NSFNET, creating more problems due to congestion and overloading. “The engineers described the situation as peeling back the layers of an onion,” explains Ellen Hoffman.

“Every time you thought you had something solid, it turned out to be only the paper layer on the top and there was another paper layer underneath it, and so you just keep going through layer after layer. You have solved a problem and all it does is reveal the next one. And so if you think of the Internet as a set of layers in terms of protocols, you could just assume that every time you hit another layer, you would have another problem, and another. It was tough!”

All of the partners were under tremendous pressure. “At times engineers were working more than a hundred hours a week,” remembers Hoffman. Knopper explains the breadth of work to be done: “We were increasing the traffic of the network, and reengineering it, and bringing on new attachments, and routing tables were exponentially growing—but what I didn't realize was that even the day-to-day stable operation of the system was having problems.” These challenges demanded everything the NSFNET team members had to give.

The last of the sixteen T3 sites was completed in the fall of 1991, over the Thanksgiving holiday, and production traffic was phased in. The new T3 backbone service, provided for the NSFNET by ANS, represented a 30-fold increase in bandwidth and took twice as long to complete as the T1 network. It now linked sixteen sites and over 3500 networks, carrying data at the equivalent of 1400 pages of single-spaced text per second. The new and improved NSFNET backbone service provided the research and education community with state-of-the-art communications networking, and connectivity to the fastest production network in the world.

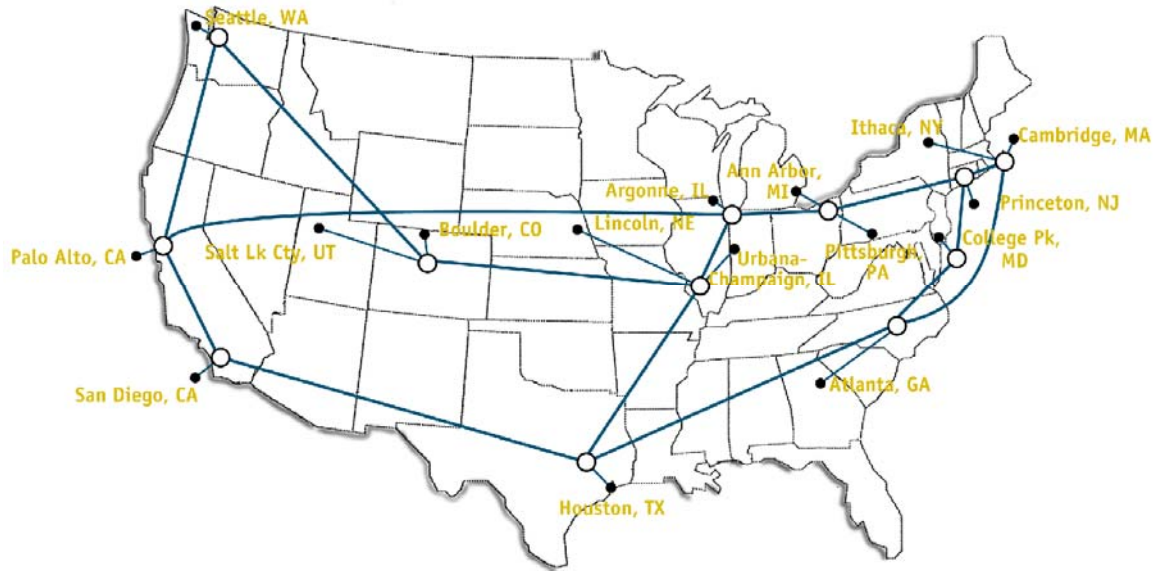


Figure 20: New T3 Backbone Service, 1992. Additional sites served by the T3 NSFNET included Cambridge MA (NEARNET) and Argonne IL (Argonne National Lab).

The introduction of a new corporate structure to the NSFNET project, and the resulting organizational complexity, created controversy among members of the research and education community, as well as other members of the Internet community. According to George Strawn, there were two main issues of concern with regard to the so-called “commercialization and privatization of the Internet.” One was what effects the NSF’s ongoing support of the NSFNET would have on the fledgling Internet service provision market; new companies such as Performance Systems International (PSI) and AlterNet charged that the NSF was unfairly competing with them. The second issue was the research and education community’s concern that “commercialization”—part of which was the perception that ANS would provide the NSFNET backbone service instead of Merit—would affect the price and quality of their connection to the NSFNET and, by extension, the Internet. They wanted to “keep it in the family,” says Strawn.

On yet another front, the regional and midlevel networks were beginning to attract commercial customers, and wanted that business for much the same reasons that ANS was created: to support themselves and the research and education community. However, they felt constrained by the NSF’s Acceptable Use Policy, which specified the nature of traffic allowed to traverse the NSFNET backbone service. Purely commercial traffic was not directly in support of research and education and was thus restricted from the NSFNET backbone. “Something had to happen to break loose the whole commercial issue,” says Ellen Hoffman.

“There had been a number of discussions about the need to be able to transmit commercial traffic, and the unhappiness with having to split your own regional network into a commercial segment and an educational segment. I have this network here in Michigan. Today it has 30 universities connected to it at the colleges, but Ford Motor wants to connect to it. If Ford Motor is not particularly interested in educational traffic, do I have to build a whole new network for Ford and create a whole new Internet for Ford? Or should I take advantage of the infrastructure that’s already there?’ That was the question we were trying to answer.”

None of these issues had touched the community before, and the NSFNET partners, including the NSF, found themselves in the midst of roiling debate. But according to those involved, the type of Internet services the NSFNET offered to the research and education community would soon be able to be obtained by the private sector. “It had to come,” says Steve Wolff,

“because it was obvious that if it didn't come in a coordinated way, it would come in a haphazard way, and the academic community would remain aloof, on the margin. That's the wrong model: multiple networks again, rather than a single Internet. There had to be commercial activity to help support networking, to help build volume on the network. That would get the cost down for everybody, including the academic community—which is what the NSF was supposed to be doing.”

Overall, the partners now realize that more steps could have been taken to reassure the research and education community that the NSFNET backbone service would continue to provide them with high-speed networking capabilities as required by the cooperative agreement with the NSF. “In retrospect, I think we should have paid more attention to better communication about our activities and objectives,” says Doug Van Houweling. Merit, IBM, and MCI believed that the formation of ANS would allow the team members to expand the network to the benefit of the research and education community. “The network we could have built with only NSF's money would not have been as robust. It would have provided connections, but it wouldn't have had the same degree of redundancy, for example,” explains Elise Gerich.

From the NSF's viewpoint, “commercial use of the network ... would further the objectives of NSFNET, by enhancing connectivity among commercial users and research and education users and by providing for enhancements to the network as a whole.”<sup>(4)</sup> However, as part of Congressional hearings related to the NSF budget, and an internal report by the Inspector General of the NSF, the project team was questioned thoroughly about its decision.

In retrospect, such growing pains were unavoidable, considering the scope of the technological achievement and importance of networking and communications infrastructure to academic (and commercial) users. Thus the upgrade of the NSFNET backbone service to T3 was not only a technological and organizational challenge of the highest order. It also precipitated a greatly-needed, though contentious, community dialogue about the evolution of the communications infrastructure that had come to mean so much to the research and education community and, increasingly, the society as a whole.



**NSFNET Networks by Country  
April 30, 1995**

Only three countries were connected by the NSFNET backbone service when the T1 network came online in 1988: the US, France, and Canada. Ten to 12 countries were added each year until 1993; in 1994, the last full year of the project, 21 new countries were added. At the end of the project, 93 countries had been announced to the NSFNET backbone service.

Country	Total Networks
Algeria	3
Argentina	27
Armenia	3
Australia	1875
Austria	408
Belarus	1
Belgium	138
Bermuda	20
Brazil	165
Bulgaria	9
Burkina Faso	2
Cameroon	1
Canada	4795
Chile	102
China	8
Colombia	5
Costa Rica	6
Croatia	31
Cyprus	25
Czech Republic	459
Denmark	48
Dominican Republic	1

Country	Total Networks
Ecuador	85
Egypt	7
Estonia	49
Fiji	1
Finland	643
France	2003
French Polynesia	1
Germany	1750
Ghana	1
Greece	105
Guam	5
Hong Kong	95
Hungary	164
Iceland	31
India	13
Indonesia	46
Ireland	168
Israel	217
Italy	506
Jamaica	16
Japan	1847
Kazakhstan	2
Kenya	1
Korea, South	476
Kuwait	8
Latvia	22
Lebanon	1
Liechtenstein	3
Lithuania	1
Luxembourg	59
Macau	1
Malaysia	6
Mexico	126
Morocco	1
Mozambique	6
Netherlands	406
New Caledonia	1
New Zealand	356
Country	Total

	Networks
Nicaragua	1
Niger	1
Norway	214
Panama	1
Peru	44
Philippines	46
Poland	131
Portugal	92
Puerto Rico	9
Romania	26
Russian Federation	405
Senegal	11
Singapore	107
Slovakia	69
Slovenia	46
South Africa	419
Spain	257
Swaziland	1
Sweden	415
Switzerland	324
Taiwan	575
Thailand	107
Tunisia	19
Turkey	97
Ukraine	60
United Arab Emirates	3
United Kingdom	1436
United States	28,470
Uruguay	1
Uzbekistan	1
Venezuela	11
Vietnam	1
Virgin Islands	4
Total	50,766

## Milestones in Technology

The NSFNET partnership reached and crossed many technology milestones over the duration of the agreement. “We defined the frontier and won it together. As a federally funded research project, few can claim such broad support and pervasive impact,” says Paul Bosco.

The most fundamental achievement was the construction of a high-speed national network service, and successfully evolving that service from T1 to T3 speeds. According to Mark Knopper, the NSFNET team's overall achievement helped to “scale networking to a huge system that covers the world. Having a single backbone in place, with a centralized administration and technical group, allowed us to do that.” The NSFNET was the first large-scale, packet-switched backbone network infrastructure in the United States, and it encouraged similar developments all over the world. Jessica Yu points out:

“The NSFNET backbone glued together the international networks—almost all traffic from abroad would transit the NSFNET. With that kind of connectivity available, other countries were prompted to build their own networks so they could get connected too. Many of them used NSFNET's three-tiered structure—backbone, regionals, campus networks—when they started to build their own networks.”

The NSFNET team played a leading role in advancing high-speed switching and routing technology. “We developed and migrated across four or five generations of routing technology through a seven-and-one-half-year period of extraordinary growth,” says Paul Bosco. “The project demanded end-to-end systems engineering with prototype deployment, but resulted in real operational experience.” Both the T1 and T3 nodes featured distributed embedded control architectures, while the T3 nodes contained powerful adapter cards that switched packets without any main system intervention, greatly increasing packet switching efficiency and making very high data throughput possible.

Many important improvements to the NSFNET backbone service went unnoticed by users. “There were many transitions,” says Hans-Werner Braun, “many times when we installed faster serial cards or upgraded other hardware and software to enrich the backbone topology. All this was invisible to users—what they saw was a network that was able to keep up with growing demands for bandwidth.”

NSFNET project partners also contributed to the development of improved Internet routing architecture and protocols. “The NSFNET backbone from day one defined a routing architecture which provided a clear demarcation between interior routing on the backbone, and exterior routing between the backbone and the attached regional and midlevel networks,” explains Yu. This separation allowed for efficient troubleshooting of network routing problems and preserved backbone routing integrity.

IBM and Merit together with Cisco Systems also undertook initial development of the Border Gateway Protocol (BGP), now used extensively on the Internet. BGP was created in true Internet style, according to Elise Gerich, “through Kirk Lockheed of Cisco and Yakov Rekhter of IBM sitting down outside an IETF meeting, working to come up with the concept for BGP and writing the specs on a napkin. It was an enormous step forward in the routing development area.”

After extensive testing, Yu and other Merit staff deployed BGP on the T1 NSFNET backbone and worked with regional engineers to update their software, so they could use BGP instead of


EGP to exchange reachability information with the NSFNET. The NSFNET backbone service was thus the first network to deploy BGP in an operational environment.

Experience gained in operating the NSFNET backbone played a major role in the development of another important routing strategy, Classless Inter-Domain Routing, or CIDR. Yu was one of the architects of CIDR, in conjunction with other members of the Internet Engineering Task Force. “Our data showed that the size of the NSFNET routing tables was doubling every ten months, and would soon become unmanageable.” she says. Yu, Elise Gerich, and Sue Hares of Merit helped convince the IETF community that a broad-based solution was needed—one that not only tackled the problem of the swiftly dwindling IP address space, but also tamed the growth of the Internet routing tables. CIDR “removes the notion of IP address classes, which reduces the size of the routing tables and enables more efficient use of IP address space,” says Yu. Technical staff from the partnership deployed CIDR on the NSFNET backbone service in 1994; CIDR is now widely used in the Internet.

In these and many other ways, the technical teams from Merit, IBM, MCI, and ANS working on the NSFNET backbone project demonstrated the effectiveness of their partnership. Their collaborative efforts to create, implement and constantly improve high-performance networking technology in the context of a rapidly changing technological environment constitute a historic chapter in the development of the worldwide Internet.

## The NSFNET Phenomenon

While the 1986 NSFNET had given the research and education community a taste of what was possible with connectivity to a high-speed backbone network, the leaps in capacity and reliability offered by the new NSFNET backbone service proved to be popular beyond even the partners' wildest dreams. The growth of the Internet worldwide, of which the construction of the NSFNET in the US played a major part, can only be described as “phenomenal.”



The statistics Merit provided on use of the NSFNET backbone are by now well-known, but still every bit as astounding: 19 billion packets per month processed in September of 1992, a growth rate of approximately 11% a month. Over 6,000 networks were connected to the NSFNET backbone service in 1992, one-third of which were outside the United States. Use of the network was also evolving; the share of network traffic devoted to electronic mail

“It was pretty much an unmanageable child,” says Ellen Hoffman. “I think back then we all had hopes that it would grow the ways that it did, but none of us believed that it would grow so fast.” Hans-Werner recalls working in an environment where “something that was brand-new and totally up-to-date in the morning was completely out of date by the afternoon. That happened pretty much on a day-by-day basis, because we moved so fast.”

Figure 11: Major NSFNET Applications by Packets, April 1995. In April 1995, WWW exceeded FTP for the first time in the number of packets traversing the NSFNET backbone service, and was only slightly behind FTP in the number of bytes


and file transfers was being overtaken by newer information storage and retrieval applications such as Gopher, Archie, and Veronica.

## Education and Outreach: To the Academic Community and Beyond

Support for the users of the network, from regional and midlevel networks to campus networks, was just as critical as the physical infrastructure, and Merit's Information Services group was tasked with giving the research and education community the background and tools they needed to use the network to its fullest. The idea of “user support” had been around since the “olden days” of campus computing systems, when “you have all this information—how do you get it so it is available to people? How do you help them use it, help them to find it?” explains Jane Caviness. Administrative information related to organization, funding and staffing of networks was also needed. Kathleen McClatchey, who as Director of Planning, Marketing and Public Relations at the University of Michigan, helped to lead the partners in these efforts from 1989 to 1991, explains that the task before them was to “try to take something at the 'bits and bytes' stage and focus on what it could actually do.” Elise Gerich puts the question this way: “What makes the NSFNET real to people?”

The Merit Information Services staff met the challenge by conceiving, designing, and providing information services for a new community of users, whose needs were very different from those of the scientists and researchers who had previously been the network's primary users. Working in parallel with BBN Systems & Technologies, which ran the first Network Information Center (NIC) for the NSFNET, the Merit Information Services staff created an array of materials designed to help new users of the NSFNET from all areas of academia, who were not as experienced with networking. Merit was also instrumental in teaching the majority of today's Internet “trainers,” who provide training materials and resources to help people learn how to use the Internet.

The Merit/NSFNET Networking Seminars, “Making Your Internet Connection Count,” brought together Internet experts, network developers and administrators, information support specialists, librarians, and teachers who needed to know the latest about networking for their institutions. Beginning in 1989, seminars were presented on-site in conjunction with NSFNET regional networks in Hilton Head, Denver, Ann Arbor, Washington, D.C./College Park, Las Vegas, San Francisco, and Orlando.



Merit Information Services staff worked successfully with many institutions to design network training on a variety of topics, from library resources to discipline-specific databases. Merit's Advanced Topic workshops educated technical personnel from the NSFNET backbone site locations about advancements in Internet engineering, network management, and routing.

Merit published the *Link Letter* newsletter to keep the community informed about changes to the backbone and the broader Internet infrastructure. A presentation called “Cruise of the Internet,” written by Merit's Laura Kelleher, introduced new and experienced Internet users worldwide to information

resources, and introduced users to tools such as telnet, FTP, e-mail, Gopher, and WAIS. A popular, computer-based Cruise was developed by Steve Burdick, Kelleher, and Mark Davis-Craig for Macintosh and Windows users.

Kathleen McClatchey and other Merit staff distributed over 1,000 copies of a video produced in conjunction with EDUCOM's Networking and Telecommunications Task Force (NTTF), "Networking to Share Ideas," containing vignettes of NSFNET users using the communications potential of the network in such diverse areas as health care, environmental research, and K-12 education. In addition, the partners produced many technical working papers and articles about networking and the research and education community in magazines and other print media.

Merit, IBM, MCI, ANS, and the NSF also maintained a high profile for the project by attending and giving presentations about the NSFNET to national, regional, and local conferences and Internet-related organizations, including the Internet Engineering Task Force (IETF), EDUCOM, FARNET, the Coalition for Networked Information (CNI), and many others. Team members also published dozens of technical papers, including Internet standards documents called RFCs (Request For Comments), that helped push the leading edge of networking technology in the Internet community. Efforts to publicize the achievements of the NSFNET backbone service began to dovetail with the larger movement toward a National Research and Education Network, or NREN. Articles in as diverse publications as *Science*, the *Wall Street Journal*, and *USA Today* helped to raise awareness of the NSFNET beyond the research and education community. Ellen Hoffman remembers the difficulty in getting people "up to speed" with the concept of networking and the Internet:

"Now you just say Internet and the press gets all excited. It wasn't true back then; invariably if you wanted to talk about the Internet, since nobody knew what it was, you would have to go through a long process of explaining. You have to remember that back then you didn't have something like the Web that made it easy to show—we had e-mail, FTP, and telnet. Those in and of themselves are very interesting applications and they do some wonderful things functionally, but they are not very showy."

This marked the beginning of an educational process for journalists and their readers and listeners that would saturate the media with stories about the "information superhighway" by 1994.

More substantively, the partners had to come up with new ways of "making the NSFNET real to people." Focusing on how users incorporated networking technology into their daily work (and increasingly, other areas of their lives), on the innovative applications being developed, and how using the NSFNET was changing the ways people related to one another seemed the best strategy. "We were trying to show the incredible power in sharing information—the idea of the collaborative work enabled by and through the network," recalls Kathleen McClatchey. "There is so much we can do with the information that we have to make people's lives better."

Today, for the growing millions of people using the Internet and commercial online services, networking is increasingly as ubiquitous as the telephone. How that happened is a testament to the technology and, more importantly, to the people using it: technology is a social construct, and often the ways in which we think about it when it is created begin to change as a result of human innovation. "I remember when we were first installing computers on campus; people installed them for one reason and ended up using them for another reason, and often they didn't really know what that other reason was going to be at the time they purchased them," says Priscilla Huston. On the ARPANET, e-mail among researchers was more of an unexpected side benefit than the stated or primary purpose of creating a secure communications system. In a similar manner, what began as primarily a research network evolved to a broader backbone network service in support of the entire research and education community, and by the early 90s was evolving yet again in the context of public discussion about a National and eventually Global Information Infrastructure. The NSFNET was one of the main catalysts for these ideas. "We

wouldn't be talking about the NII without the NSFNET," says Ellen Hoffman. That particular achievement is one of the NSFNET's proudest.

### **Realizing the Vision: Lessons and Insights from the Project Team**

NSFNET team members agree on the primary factors that led to the success of this project: vision, flexible yet firm management, and commitment from individual staff at every level of the partnership. Over and over again, staff from Merit, IBM, MCI, ANS, and the NSF emphasized how important it was to have a clear objective guiding their way, and that the foundation for that goal—building a national backbone network—was a vision of what it could mean to people using it. Doug Van Houweling puts it this way:

“We all shared a sense of what this could be, and when the going got tough on getting the circuits to work or when people questioned our competence or motives, we fell back on our conviction 'This is Important, This is Significant.' A lot of people look for an opportunity to have a real impact on the world. When the chance arrives, it provides powerful motivation.”

Clear objectives, however, didn't mean a rigid adherence to what had been done before, or to only one notion of how to do it. In a dynamic environment where, as Hans-Werner Braun noted, expectations and events would change on a day-to-day basis, flexibility was key. “Disguised limits, no preconceived limitations,” is how Mathew Dovens puts it. “This project was set up to open new horizons, and I think throughout the seven years we never limited ourselves.” Flexibility was balanced with firm guidance, and by all accounts Merit and the NSF provided just the right mix of incentives. Harvey Fraser notes particularly the NSF's contribution:

“Steve [Wolff] had very high goals, and did not interfere with implementation, but many times there were choices to be made which may or may not match the most recent proposal that's been funded. In the process of exploring those choices, it always came out that what he wanted was technological leadership and excellence of operation, and however that got manifested in a given situation, it didn't matter. If it meant that there was a need for the government to invest more money, buy more circuits, build something else, whatever it was, fine. He certainly provided a very steady hand on the tiller.”

Other team members singled out Eric Aupperle as having managed the building of the NSFNET backbone well: “Eric was always a gentleman—he was forceful in expressing his views, but was always ready to find a compromise.”

Perhaps the most important contributor to the success of the NSFNET project, however, was the commitment among the partners—Merit, IBM, MCI, ANS, and the NSF—at every level, to make it come together.

All of the NSFNET team stressed the importance of “buy-in” to the project on the executive, managerial, and technical staff levels of the partnership. Part of this was the individual commitment to the overall “vision” of the NSFNET. But a willingness on the part of each person to commit to the job before them was necessary as well, even if it meant long hours, excruciatingly difficult technical problems, or criticism from the community.

Another aspect of this commitment showed in the partners' willingness to collaborate closely throughout the length of the project, through a series of technical and administrative meetings that took place on a weekly, monthly and quarterly basis. Finally, in the back of their minds, the NSFNET team members knew that if they didn't succeed, there was no other alternative for the

community. According to the NSF and Merit, they recognized that as the NSFNET was at first the only backbone network, they couldn't let the community down and were obligated to put the very best network in place. Mark Knopper says:

“This was it: there was no other backbone. There was no competition. The user community was there, demand for the bandwidth brought in the traffic, the regionals aggregated it, we had to carry it. That forced us to look ahead, time and again, because no one else was in a position to do it.”

Immense personal satisfaction came from working as a team where commitment, both to the project and to a larger vision, flexibility, and skilled management were part of the environment around them. “To make it happen and really make a difference, and to drive the whole thing forward,” says Hans-Werner, was the most rewarding. Mark Knopper speaks for them all: “I was honored by working with those people.”

What lessons can be learned from this successful partnership between the research and education community, industry, and government? Could the nation do it again? “I think it was unique, but it doesn't have to be,” says Elise Gerich. Mark Knopper emphasizes the importance of NSF funding as well as industry cost-sharing; according to him and others, “leveraging was a really, really big win” for the project. Steve Wolff thinks that similar initiatives would benefit from “finding out when the interest of the academic community and industry are similar, and forming a partnership with government participating at the margin. It's a great model,” he says.

Also important was the way in which the NSF and Merit could bring different players together: “A central element of the success was the extremely productive interaction between industry, academia, and government,” says Paul Bosco. Walter Wiebe notes that “being part of the team that pulled it together will be a source of satisfaction to me all my life.” According to Priscilla Huston, “the NSFNET project has involved a lot of different parties, not just the NSF, but international and industry players. Merit's support of that was an important part of the award.”

Last but not least, the way in which the funding was applied, the overall construction of the NSFNET program—of which the backbone service was the key piece—was done in a way that would reap the maximum benefits from infrastructure building. “You really have to be building at both levels at the same time,” explains Ellen Hoffman.

“In creating a ubiquitous network, you have to do both. You want it to be widespread—first of all the NSF built a backbone with lots of bandwidth, a playground for people to work with. At the same time they were doing other things to start making it widely available, like the Connections program; almost every college in this country started off with an NSF grant for its first connection. That is the other part of the ubiquity. And they didn't want to do only one or only the other, because they recognized that if they hadn't built the backbone we wouldn't have resources like the Web.”

All of these factors combined, say the partners, helped to make the NSFNET a success, to build the very best national backbone network of its kind at the time. Says Al Weis: “What the NSF accomplished was outstanding. They took a little bit of seed money and created the right dynamics which, in essence, sparked the Internet, the global Internet.”

### **III. Transition to the Future**

#### **A New Architecture for the NSFNET**

By the summer of 1991, as the NSFNET partners continued to work on the T3 upgrade, the NSF knew that the five-year cooperative agreement with Merit would be ending in 1992. Around this time, according to George Strawn, NCRI began thinking and planning in earnest about what to do next. It seemed clear that in the context of the debate about commercialization, continuing on the same course of support for the NSFNET was not an option; nor was the NSF particularly interested in doing so. Steve Wolff saw the actions of the NSF with the NSFNET as similar to the NSF's support of campus computing centers in the early sixties and seventies, which was phased out after computing products and services were readily available from the marketplace. The NSF knew in both cases when to initiate high-technology initiatives, and when to get out. Wolff explains the NSF's philosophy in this way:

“I think it's a strength. You're subject to criticism, because people do get hurt in the process; there's no question about that. On the other hand, NSF recognizes limitations and only has so much money. If you don't stop doing old things, then you can't start any new things. And when something gets to the point that it becomes a commodity product, there is no reason for NSF to be supporting it.”

Internet service providers were springing up all over the country, from local dial-up providers to larger companies providing T1 and eventually T3 service, and there were now a number of vendors offering TCP/IP networking products and services. During 1992, the National Science Board authorized an extension of Merit's cooperative agreement for eighteen months beyond the October 1992 expiration date in order for NCRI to develop the solicitation. Recognizing that the best way to solve some of the thorny issues precipitated by the pending termination of the cooperative agreement—and the commercialization question—was to involve the community in its plans more directly, in the early summer of 1992 the NSF distributed a draft of its solicitation for a new architecture to the public. The design for the new architecture was based on concepts developed in a paper titled “NSF Implementation Plan for Interim NREN” by Bob Aiken of the NSF; Hans-Werner Braun, then at the San Diego Supercomputer Center; and Peter Ford, on assignment to the NSF from Los Alamos National Lab.(5)

The draft solicitation was written by George Strawn, then NSFNET Program Director under Wolff at NCRI. “Steve said, ‘This is really complicated, and it will be really important for us to get it right. I think we need to put out a public draft,’” remembers Strawn. NSF received over 240 pages of comments on the draft; Peter Ford and NSF's Don Mitchell helped Strawn incorporate the comments into the final plans. Finally, in the spring of 1993 the NSF released Solicitation 93-52, *Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for NSFNET and the NREN Program*.

The solicitation contained four parts: a very high-speed backbone, or vBNS; Network Access Points; a Routing Arbiter; and Regional Network Provider awards. The vBNS, linking five NSF supercomputer centers and operating at minimum speeds of OC-3 (155 Mbps), would offer a high-speed, high-bandwidth virtual infrastructure for networking research, including advanced applications. Network Access Points, or NAPs, would act as interconnection points for commercial Internet service providers. The Routing Arbiter would manage the ever-growing routing tables and databases for the providers connecting at the NAPs. Finally, regional and midlevel networks would continue to receive NSF funding, phased out over a four-year period, to



support their connections to the Internet. They would use the money to pay commercial Internet service providers, who were then required to connect to NAPs. In this way, the NSF hoped, connectivity for the research and education community to the rest of the Internet would be maintained.

During 1993, as the responses to the solicitation came in to the NSF, the partners continued to maintain the NSFNET backbone's stable networking environment. Various hardware and software improvements to the backbone were added, and a major router upgrade doubled the speed of packets switched on the network. Merit also introduced its new Policy Routing Database, or PRDB, an expansion and improvement of routing procedures used on the NSFNET backbone service since its inception. At the same time, Merit prepared and eventually won its proposal for a Routing Arbiter award. All of the winners of awards for the new architecture were announced throughout 1994. Merit and the Information Sciences Institute (ISI) at the University of Southern California together form the Routing Arbiter team.

The vBNS award was given to MCI; Network Access Point Manager awards were given to Sprint, for a New York NAP; MFS Datanet for a Washington D.C. NAP; and Bellcore, which manages two NAPs: the Chicago NAP in conjunction with Ameritech, and the California NAP in conjunction with Pacific Bell. Awards to the regional and midlevel networks for interregional connectivity were given to seventeen regional and midlevel networks.

In May of 1994, Merit's cooperative agreement was extended one last time through April of 1995. From 1994 through the first half of 1995, Merit and NCRI worked closely with the regional and midlevel networks to ensure a smooth transition, helping them to disconnect from the NSFNET backbone service: "Merit's award was winding down, but it was their responsibility to get people moved over and they did a really good job pulling people together," says Huston. The PRDB was also succeeded by the Routing Arbiter Database, a global registry of routing and networking information that forms an important part of the new Internet Routing Registry.

MCI began to operate the vBNS as well as provide networking services to many of the NSFNET's regional and midlevel networks through internetMCI. IBM formed Advantis, the US-based part of the IBM Global Network, with many of the personnel from the NSFNET project. The NSFNET partners also continued to provide T3 service to the NSF-sponsored supercomputer centers until the vBNS came online. On April 30, 1995, the NSFNET backbone service was decommissioned, and an era ended—in triumph.

## **Conclusion**

The NSFNET backbone service will be remembered not only for achieving the goal of providing networking connectivity to the research and education community, but also for realizing some of the larger hopes of the partners: that networking would become a ubiquitous part of everyday life for more and more people, and that it would create new markets for products and services in networking and communications, thus promoting technology transfer.

Merit, IBM, MCI, ANS, and the National Science Foundation knew that their immediate task was to improve on the previous NSFNET backbone, to push the technology past what had ever been attempted before, and to implement it on a national scale. In doing so, they hoped to place a powerful tool in the hands of the research and education community, and spur innovation in the development and use of communications infrastructure and applications. The partners were successful on both counts: the NSFNET backbone service connected most of the higher research and education community in the US to a robust and reliable high-speed networking

communications infrastructure. It also played a singular role in creating an Internet “industry,” a dynamic new marketplace for communications technologies. As the NSFNET spurred the growth of the Internet, moreover, use of networking technology expanded outward from scientists and educators in academia and industry to encompass individuals in all kinds of enterprises and organizations, both private and public, all in a few short years.

The NSFNET program will continue to evolve, to reflect the changing needs of the research and education community as well as the revolution in communications infrastructure it had such a large hand in creating. Today's new architecture does not represent a “new” NSFNET. Rather, the NSF sought awardees to provide the main structural elements of the next generation of high-speed networking in the US as well as internationally. To accomplish these goals, the NSFNET program is moving in two directions: to provide support for the research and education community by ensuring the availability of services, resources and tools necessary to keep the larger Internet interconnected, and to continue to push the envelope of networking technology with projects such as the vBNS—all in the context of the burgeoning market for commodity providers of Internet services and applications.

The success of the NSFNET project demonstrates the potential of partnerships between academia, business, and government in high-technology initiatives. The keys to its success seem to lie in the combination of vision, technical achievement, commitment, and luck that marked the project throughout its seven and one-half years of operation.

The partners began with a vision of what the NSFNET backbone service could mean to the research and education community, enabling them to work collaboratively, no matter where they were physically located, and share the results of their intellectual productivity in seconds. This vision was supported by commitment to achieving the project goals at all levels of the organizations involved: by all of the programming and engineering, information and user services, managerial and executive staff at Merit, IBM, MCI, ANS, and the National Science Foundation. An important part of that vision was the philosophy of building infrastructure at two levels simultaneously: constructing the NSFNET backbone service at the top and providing for regional network connectivity while empowering individuals at the base by funding campus connections to the NSFNET. Another important part was to expand the original notions of research networking to include education, following the example set by modern research universities that attempt to combine teaching expertise with research activity.

The partners reached the technology frontiers of networking and communications and crossed them—from building a nationwide T1 backbone service to reengineering, expanding and upgrading a nationwide T3 backbone service that at its peak transferred over 86 billion packets per month at 45 Mbps speeds—the equivalent of the contents of the Library of Congress every two weeks. At the same time, the NSFNET partners also had to respond to daunting challenges, both technical and otherwise, engendered by building networking infrastructure on a national scale: as the NSFNET (and the Internet) exploded in popularity, a complete upgrade to T3 speeds was required to support increased use, which in turn provoked questions of making the technology self-supporting by commercializing the technology. The transition from the NSFNET backbone service to a new architecture incorporating commercial Internet service providers resulted from the partners' creation of ANS and involvement in community discussion of issues of “commercialization and privatization” of the Internet. All of these elements accounted for the project's success, together with a bit of good fortune: “I don't think there's any substitute for luck,” Doug Van Houweling says. We can learn much from the example set by the NSFNET.

Since the earliest days of the telegraph and the telephone, history tells us that the arrival of each new communications medium has been accompanied by grandiose claims of its potential benefits to society. In order to take advantage of the exciting opportunities afforded by today's technology, it is imperative that policy makers examine the development of the NSFNET and the Internet. We are still far away from a truly open, interoperable, and ubiquitous global information infrastructure accessible to all, "from everyone in every place to everyone in every other place, a system as universal and as extensive as the highway system of the country which extends from every man's door to every other man's door," in the words of Theodore Vail, president of AT&T in 1907. However, the Internet has brought us a giant step closer to realizing the promise of high-speed networking, one of the most revolutionary communications technologies ever created. As part of this phenomenon, the NSFNET backbone service provided a model for future partnerships as well as a legacy of technology for the world.

(1) Dennis Jennings, Laurence Landweber, Ira Fuchs, and David Farber, "Computer Networking for Scientists," *Science* 23 (February 1986).

(2) EDUCOM, Networking and Telecommunications Task Force, "A National Higher Education Network: Issues and Opportunities." May 1987.

(3) Office of the Inspector General, "Review of NSFNET." 1993: National Science Foundation 93-01, p. 38.

(4) Office of the Inspector General, "Review of NSFNET." 1993: National Science Foundation 93-01, p.68.

(5) B. Aiken (NSF), H.-W. Braun (SDSC), and P. Ford (LANL); ed. K. Claffy (SDSC), "NSF Implementation Plan for Interim NREN" (May 1992), *Journal of High Speed Networking*, Vol. 2, No. 1, 1993.

## ***Acknowledgements***

This report was researched and written by Karen D. Frazer.

Merit Network, Inc. would like to thank all those who contributed to the report: Neal Lane, Ellen Hoffman, Mike Roberts, Jane Caviness, Priscilla Huston, Doug Van Houweling, Dale Johnson, Harvey Fraser, Mark Knopper, Elise Gerich, Al Weis, Eric Aupperle, Don Mitchell, Steve Wolff, Walter Wiebe, Mathew Dovens, Jack Drescher, Paul Bosco, Larry Bouman, Hans-Werner Braun, Dan van Bellegham, Bob Mazza, Rick Boivie, Alan Baratz, Kathleen McClatchey, Jessica Yu, and George Strawn.

Many dedicated people helped plan, build, operate, and promote the NSFNET backbone service. It is hoped that this report will serve to recognize the many team members, named and unnamed, who helped to build the NSFNET.