NSF Hands Over the Internet ...

Next month marks the privatization of NSFnet and creation of a new research link, while its four supercomputing centers come under scrutiny

Cheered on by the new Republican-dominated Congress, Washington is awash with talk about turning over federal programs to the private sector. But the National Science Foundation (NSF) is poised to do more than preach: Next month it will turn rhetoric into reality by privatizing a service used by virtually every scientist in the country.

The change affects the Internet, the booming web of electronic networks that serves more than 20 million users in 60-plus countries. When the conversion is complete, the federal government will no longer be directly supporting this electronic highway. The immediate impact for scientists is likely to be slightly higher charges. But government and industry officials are betting that long-term costs will drop, and they expect a new NSF-funded system for researchers will lead to the next great advance in computer networking.

More than 2 years in the making, the transition to the next electronic era has not been easy. It also offers sobering lessons for those eager to see the private sector handle what had once been a government responsibility. When Congress and the White House agreed in 1992 to privatize NSFnet, which forms the backbone of the Internet, they left the details to foundation officials, who had already begun to explore the idea. "We thought we'd be able to do this in 18 months, or maybe sooner," says Jane Caviness, who directed the University of Delaware's computer center and now serves as temporary director of NSF's networking and communications research division, which is in charge of the changeover. "The assumption was you could just spin this off," adds Fred Weingarten of the Computing Research Association.

That assumption was wrong. For example, NSF split its network operations into several parts, each of which is being privatized separately, to avoid creating a commercial monopoly. But this approach proved time-consuming. So did debugging new technology designed to cope with the increasing flow of traffic on the network and resolving a legal challenge by a losing bidder for the new high-speed network. "There was pressure to get [operations] out of NSF without really thinking things through," says one university researcher. "And the politics did not line up with the state of technology."

Now, however, NSF officials say they have overcome the bulk of these obstacles. "There are no more showstoppers," says Caviness. Adds a congressional staffer familiar with the effort, "We're not hearing there are any serious problems."

Four not-so-easy pieces

In privatizing NSFnet, NSF officials divided the work into four parts. The first involved developing access points for the various regional networks. These are being developed by Sprint, Bellcore (and its subsidiaries Pacific Bell and Ameritech), and MFS Datanet. The second was selecting a system traffic cop to treat network services providers equitably and to simplify procedures in the unruly Internet. The third was strengthening the regional networks that connect large numbers of users, and the last was developing an advanced high-speed research network.

The network access points—the hubs where regional networks tap into each other as well as into national and international networks—have been the hardest nut to crack, NSF and industry officials say. The hardware and software developed to handle the immense traffic at these points was untested, explains Steve Wolff, who was in charge of the transition at NSF until he left in January to join private industry. "It turned out to be not quite ready for prime time."

As a result, the companies had to come up with an interim system while the more advanced one is still being tested. Mike Roberts, vice president for networking at Educom, a Washington-based organization which represents dozens of universities, blames those companies for the delay. "There was so much ego," he says. "They said, 'We know better.' Finally, after this purist nerd stuff, we're working together." Rick Hronicek, Pacific Bell's executive director of advanced computing, does not deny there were problems, but he says that industry was simply doing what NSF had requested.

NSF managers also had to abandon plans for a comprehensive information service designed to help users pick from among private network services. The foundation canceled a $4.8 million agreement with General Atomics because of a dispute over the company's management of the project. And the flood of networking information available commercially in recent years made the service largely unnecessary, NSF managers say.

Meanwhile, MCI Corp. is finishing work, delayed a year because of a challenge by Sprint, on a new system for researchers. The very high-speed Backbone Network Services (vBNS), connecting NSF's four supercomputer centers, will be able to handle far more data, with greater efficiency, than NSFnet. The idea is to let vBNS grow like NSFnet did—starting with a small and select group of researchers. Unlike current NSFnet traffic, the first vBNS users will be those with very high-speed computing applications that have been

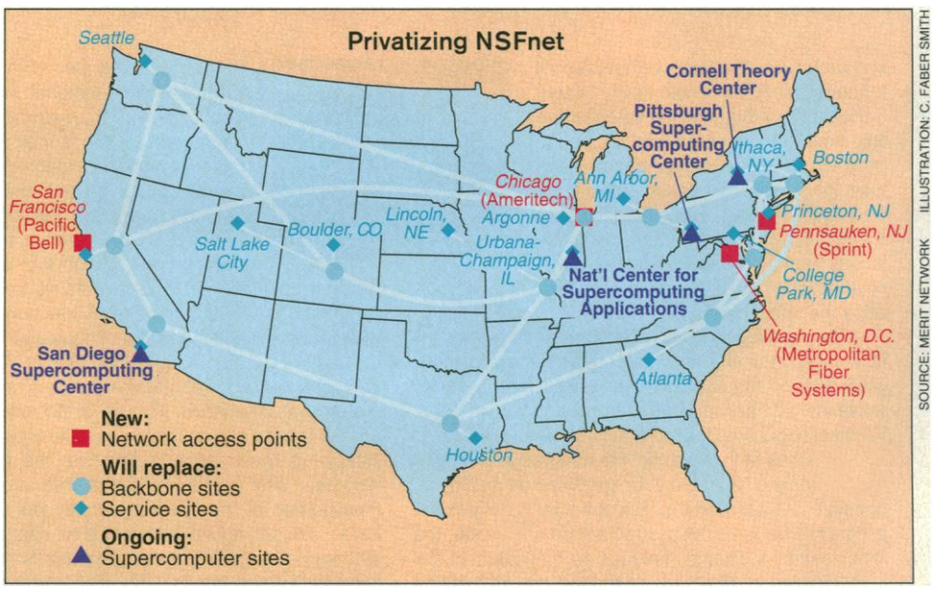

Shifting landscape. Network access points will replace backbone and nearby service sites.

peer-reviewed. But major universities are already excited by the advent of the system; a group of senior university officials met earlier this month in San Francisco to discuss the subject. "It's going to be a hit; it could trigger a whole new level of activity," says Wolff. The system is slated to be operating by 1 April. "The current NSFnet system is just not going to support the kinds of high speeds needed," says Lawrence Rowe, a computer scientist at the University of California, Berkeley. "For videoconferencing and sharing documents, we need a lot more bandwidth."

Fear of higher fees

Although the future is bright, the present is clouded with worry about costs. Universities and colleges pay from $10,000 to $60,000 a year for access that regional networks provide to the backbone subsidized by NSF. Now those regional networks will have to pay commercial providers for access to the backbone, and the cost of using the networks inevitably will be passed on to the users. How much, however, is hard to forecast.

Marjory Blumenthal, a National Research Council staffer who directed a study last year on the future of information systems, says costs could go up on average from 10% to 100%, depending in part on one's distance from an access point. The study, *Realizing the Information Future*, warns that small colleges, public libraries, and schools with small budgets will feel the pain more sharply than wealthier institutions will.

At the same time, individual researchers likely will not be billed directly for the higher costs, say university and NSF officials. Instead, the charges will be added to university overhead rates. And NSF intends to cushion the blow by subsidizing the regional nets, in diminishing amounts, over the next 4 years. "The change in NSFnet is causing more apprehension than circumstances appear to warrant," the study notes.

All in all, most observers see the changes in a positive light. "The scientific community will be better off in the new environment because the net will be operated more efficiently by the private sector—and therefore at a lower cost," says Representative Rick Boucher (D–VA), who helped frame the issue in a bill to privatize NSFnet that was introduced in 1992 but was never enacted. And Educom's Roberts thinks the changes will be taken in stride by most scientists and universities. "The research community is pretty comfortable with this," says Roberts. "After all, it's been involved with this technology now for 10 years."

Despite the lingering uncertainty about the Internet's future, everyone agrees on one point: Researchers will find new uses for networks. And those uses will increase traffic and run up the tab for their institutions.

–Andrew Lawler

U.S. SUPERCOMPUTING

... Weighs Future of Computer Centers

Charles Peskin isn't a medical doctor, but for the past 25 years he's been operating on the human heart. An applied mathematician at the Courant Institute of Mathematical Sciences at New York University, Peskin does his "surgery" via computer, using a three-dimensional model to describe blood flow through this vital organ. The model, developed with colleague David McQueen, is so complex that a single heartbeat takes 150 hours to run on a Cray Research C90 supercomputer. And the machine isn't even located in the Big Apple—it's nearly 400 miles away at the Pittsburgh Supercomputing Center.

Peskin's work, which won two national computing awards last year, is helping scientists to understand heart disease and companies to design better valves. It's made possible by a program at the National Science Foundation (NSF) that has provided 15,000 researchers at 200 universities with access to supercomputers over the past decade. In 1985, when supercomputing was in its infancy, NSF established five centers, equipped them with state-of-the-art machines, and linked them together with a high-speed network that later formed the basis of the Internet. The centers make computing time available to researchers like Peskin, free of charge, on the basis of a peer review of proposals. This year NSF will spend almost $60 million on cooperative agreements with Pittsburgh and the three remaining centers—the Cornell Theory Center, the National Center for Supercomputing Applications (NCSA) at the University of Illinois, and the San Diego Supercomputing Center. (In 1989 NSF decided to end support for a center at Princeton University.)

But that arrangement could be changing. Last fall the National Science Board, NSF's governing body, turned down an NSF-backed proposal to renew the centers' contracts for 5 years. Instead, the board extended the existing agreements for 2 years and directed NSF officials to conduct a thorough review of the centers' operations. One is needed, they say, because scientists can satisfy their supercomputing needs at other sites and because a tight NSF budget compels them to scrutinize every large program. These moves have sent a tremor through the centers and their growing band of users, who fear that researchers' access to supercomputers may be reduced.

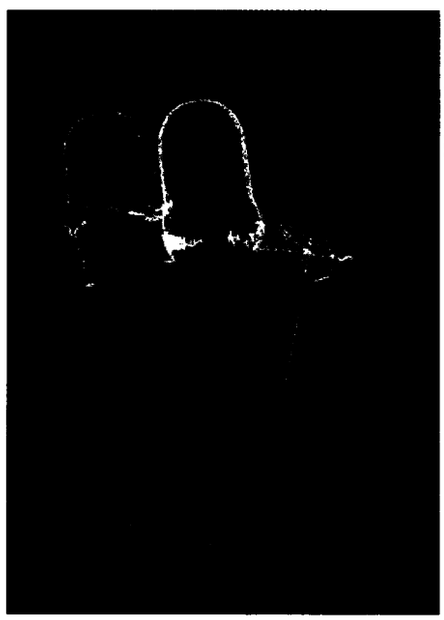

PESKIN AND MCQUEEN

In the flow. A cross section of Peskin and McQueen's three-dimensional heart model shows computed blood flow during ventricular ejection.

A task force,* chaired by Edward Hayes, vice president for research at Ohio State University, convened in January in the hope of delivering its report to the board next November. The eight-member panel will analyze "all possible alternatives," says Paul Young, head of NSF's computer and information science and engineering directorate—including giving vouchers to scientists to spend on supercomputing as they choose, requiring them to obtain outside funding for computing, holding an open competition for future center sites, or terminating the program.

Whatever the panel decides, its recommendations are likely to alter an arrangement that, in the words of board member Thomas Day, president of San Diego State University, has been "incredibly successful" for a decade. Set up to provide U.S. academic scientists with access to supercomputers, the centers have moved beyond what are known as "flop shops"—sources of raw computing power—to becoming service centers that offer technical support, programs like Mosaic (developed at the Illinois center), digital libraries, visualization, communications, and education.

* The task force members are: Arden L. Bement, Purdue University; John Hennessy, Stanford University; John Ingram, Schlumberger Research Laboratories, Austin, Texas; Peter A. Kollman, University of California, San Francisco; Mary K. Vernon, University of Wisconsin; Andrew B. White Jr., Los Alamos National Laboratory; and William A. Wulf, University of Virginia.