

Retiring the NSFNET Backbone Service: Chronicling the End of an Era

By Susan R. Harris, Ph.D., and Elise Gerich

Reprinted with permission from "ConneXions", Vol. 10, No. 4, April 1996.

April 30, 1996, marks the one-year anniversary of the final dismantling of the venerable NSFNET Backbone Service. After more than a year of planning, reconfiguration, shutdowns, and transitions, the U.S. Internet had completed its move to a new architecture composed of multiple backbones, linked at the new interexchange points.

The midnight NSFNET shutdown went remarkably smoothly, as did most of the events leading up to the final phaseout. This article looks back on the timelines, dependencies, delays, emergencies, and successes that marked the final year of the NSFNET. We begin by taking a brief look at the history of what was the world's largest and fastest network for research and education.

A Brief History of the NSFNET

The National Science Foundation inherited the responsibility for nurturing the U.S. Internet from the Advanced Research Projects Agency (ARPA). From its inception in 1985-1986, the NSFNET program laid the foundation of the U.S. Internet and was the main catalyst for the explosion in computer networking around the world that followed.[1] The first NSFNET, a 56Kbps backbone based on LSI-11 Fuzzball routers, went into production in 1985 and linked the six nationally funded supercomputer centers (the five NSF centers and the National Center for Atmospheric Research). Soon after the network's inception, the need for more advanced networking technology was indicated when rapid growth in traffic precipitated serious network congestion. In 1987, NSF issued a competitive solicitation for provision of a new, faster network service. The new service would provide a network backbone to link the six supercomputer centers and seven mid-level networks. The mid-level networks would in turn connect campuses and research organizations around the country, creating a three-tiered network architecture that remained in place until the end of the NSFNET backbone service.

In fall 1987, NSF selected Merit Network, Inc., and its partners MCI, IBM, and the State of Michigan to manage and re-engineer the new backbone service. Eight months after the NSF award, the NSFNET partnership delivered a new T1 backbone network that connected 13 sites: Merit, NCAR, BARRNet, MIDnet, Westnet, NorthWestNet, SESQUINET, SURAnet, and the NSF supercomputer centers. Two additional regional networks, NYSERNet and JVNCnet, were also served by the backbone, because each was collocated at a supercomputer center. Each of the 13 backbone nodes, known as Nodal Switching Subsystems, was composed of nine IBM RTs linked by two token rings with an Ethernet interface to attached networks. There were 14 T1s connecting the sites, on which a virtual topology was constructed. Each virtual path represented one-third T1 to the site.

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.[2] With this upgrade, the NSFNET's physical topology equalled its virtual topology. By then, 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.[3]

Upgrade to T3

To handle the increase in traffic, Merit and its partners introduced a plan to upgrade the backbone network service to T3. The NSF also wanted to add a number of new backbone nodes, and asked Merit to prepare proposals for the added cost of new nodes at T1 and T3 speeds, while the NSF issued a solicitation to the community for those interested in becoming new NSFNET sites. It was eventually decided by the NSF that the partners would increase the total number of backbone nodes on the NSFNET from 13 to 16, all running at 45 Mbps.[4] Additional sites served by the T3 NSFNET backbone service would include Cambridge MA (NEARNET), Chicago's Argonne National Lab, and Atlanta GA (SURAnet).

In late May 1990, Merit's cooperative agreement with NSF was modified to cover the additional work. By the end of the year, Merit, SDSC, and NCSA were connected to an early T3 service and began testing the new T3 routers with real traffic. In addition, a new T3 research and test network was implemented to parallel the existing T1 test facility.

Important architecture and equipment changes came with the new T3 network. The core backbone equipment was moved from the universities and supercomputer sites to MCI's points-of-presence (POPs), and the RTs were replaced with RS/6000s and a card-to-card forwarding architecture. Many of the techniques introduced in the T3 RS/6000 routers have since been adopted by commercial router vendors.

As the backbone network service was growing in complexity and was re-engineered, increasing focus and resources were needed to keep pace with more complex technical, business, and policy environments. To meet these organizational challenges, ANS was created and announced in September 1990. ANS began to provide service for NSFNET as a subcontractor to Merit, with IBM, MCI, and others continuing to infuse new technology to develop the infrastructure.

During 1991, a year of refining the new backbone technology, the T1 and T3 networks existed in parallel. Difficulties in tuning the new technology prevented the network from being moved to full production status until late in the year, when all sixteen backbone sites comprising the NSFNET service were connected to the new ANSnet national T3 infrastructure. With expansion work completed and improved performance validated, several sites began using the T3 for their primary traffic path by November 1991. A final round of testing in mid-December set the stage for moving the remaining NSFNET traffic to the new backbone service in early 1992. The network now exceeded the T1 structure in stability by a factor of ten, with fewer outages and errors in all categories.

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 and commercialization of the U.S. Internet. 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 growing 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 NSF to develop a follow-on solicitation for national networking, one that would accommodate the growing role of commercial providers and allow NSF to step back from actually operating a network to concentrate on supporting leading-edge research initiatives. NSF published a draft solicitation for community comment in 1992, and a new solicitation was issued in May 1993.

Early in 1994, awards for building the new architecture were given to Merit and USC's Information Science Institute for the Routing Arbiter service, to MCI for the vBNS, and to three providers for the Network Access Points: Sprint, MFS Datanet, and Bellcore, representing Ameritech and PacBell. NSF also awarded Merit a transition extension that began in May 1994 and lasted until April 1995, when the NSFNET backbone service would be retired and all connections would be switched to a new service.

Deadlines and Commitments

Moving the U.S. Internet to a new architecture in the months between the 1994 awards and the April 30, 1995 termination date was a frightening challenge for the regional networks, the ISPs, and the NSFNET partnership. Before the backbone could be decommissioned, four main tasks had to be accomplished by the networking community:

- Establish the Network Access Points (NAPs) and move them to production status.
- Attach to the NAPs the NSFNET and the ISPs that provided service to the regionals.
-
- Develop the RA Service by placing Route Servers at the NAPs and setting up a routing registry.

Move the regionals off the NSFNET and attach them to networks operated by ISPs.

According to NSF's ambitious transition schedule, the new NAPs would be available by August 15, 1994. The NSFNET backbone service would then attach to the NAPs, with all current attachments to the NSFNET remaining in place. The ISPs would then begin to attach to the NAPs, and regional networks that attached to NSFNET would begin to establish connections to the ISPs. By October 31, the regionals would cut over all traffic to the ISPs and disconnect their attachments to the NSFNET. Only the supercomputer centers would remain attached to the NSFNET. The vBNS would be deployed by January 1, 1995, and attached to the NAPs by February 1, 1995.

As it turned out, all of these actions were delayed, and revised deadlines established.

Establishing the NAPs

The first Network Access Point to go into production was the Washington, D.C. NAP (MAE-East, the Metropolitan Area Ethernet). MFS had been operating MAE-East since 1992, and MAE-East had served as a model for the NAPs as defined in NSF's solicitation. In fall 1994, MAE-East was upgraded from a 10Mbps Ethernet to FDDI; internetMCI and SprintLink, which had already attached to the MFS facility, upgraded their connections to FDDI, as did the NSFNET.

The Sprint NAP, a bridged FDDI/Ethernet hybrid, was up and running by the end of the summer; ANSnet/NSFNET, SprintLink, and internetMCI attached to it in September. The Sprint and Washington, D.C. NAPs began to carry much of the traffic for the U.S. Internet once networks began to move off NSFNET in November 1994, because the PacBell and Ameritech ATM NAPs were still being deployed and went into production several months later. Both facilities were physically in place by October 1994, but problems with ADSU performance and a concern with ATM switch buffer sizes led to a lack of confidence in the ability of the ATM NAPs to sustain the traffic load.

As a result, both PacBell and Ameritech decided to deploy interim configurations, and put FDDI LANs into production in March 1995. Some ISP routers on the FDDIs at these contingency NAPs were also connected to DS3 ports on the ATM switch, so they could pass traffic across the FDDI while still transmitting to ATM-connected peers. As of January 1996, this infrastructure is still in place at the PacBell and Ameritech NAPs.

Deploying the Route Servers

The Routing Arbiter service has two main components: the Route Servers, SPARC 20s deployed at the NAPs, and the Routing Arbiter Database, successor to the Policy Routing Database used to configure the NSFNET backbone service.

In November 1994, primary and backup Route Servers were shipped from ISI to each of the NAPs. Once the necessary data circuits, front-end systems, controllers, ATM switches, and FDDI bridges were installed and tested, addressing schemes worked out, security procedures implemented, and 24/7 network monitoring in place, the Routing Arbiter team began to set up peering sessions with customer routers at the NAPs. Out-of-band access -- a prerequisite for declaring the Route Servers fully in production -- became available several months later.

By April 1995, the Route Servers were peering with more than a dozen providers at the Sprint and Washington, D.C. NAPs. In July, production RA services were announced at the Sprint NAP, and announcements for the other NAPs soon followed. At each exchange point, the Route Servers began importing and exporting routes to numerous ISPs. The ISPs maintained sessions with other peers as well as the Route Servers, comparing the routing information from both sessions for consistency.

NACRs and the PRDB: the Long Goodbye

Merit originally planned a December 1994 retirement for the Policy Routing Database (PRDB), which had been used to configure the NSFNET's backbone routers since 1989. The PRDB would be replaced by the Routing Arbiter Database, which would then become part of the Internet Routing Registry (IRR) along with the RIPE NCC, MCI, ANS, and CA*net registries. The IRR would be an important global resource--a public repository of announced routes and routing policy in a common format, so that ISPs could use the information stored in any and all registries to configure their backbone routers, analyze routing policy, and build tools to help in these efforts.

The PRDB was established to maintain information about what were considered legitimate destination announcements from the various regionals. The primary goal of maintaining this information was to prevent routing loops. When BGP replaced EGP as the inter-domain routing protocol in 1994, suppression of routing loops no longer had to be so administratively controlled. The information in the PRDB was then mainly used to record routing policies such as path preferences and to generate the backbone configuration files.

NSF's follow-on solicitation for the new architecture specified a continuation of the function that the PRDB played in the T1/T3 NSFNET. The goal was to record global routing policy information based on each Autonomous System's policy. RIPE had pioneered this work in the European arena, and the data exchange format described in RIPE-181 (RFC 1786) was adopted as the "standard" for Internet Routing Registries.[5] The RADB adheres to this model.

The challenge was to establish and populate the RADB before the retirement of the NSFNET and the PRDB. By summer 1994, the RIPE NCC registry had been in production for two years, and CA*net and internetMCI were creating routing registries to support their customers. ANSnet would continue to use the PRDB until the RADB was established. But the dilemma was how to transition from NSFNET-centric information to the AS-specific information needed for the RADB, while continuing to provide a stable router configuration environment for the NSFNET service.

Merit's December target date for retiring the PRDB was based on the assumption that the regionals would be off the backbone by October 31. When it became clear that they weren't going to make that deadline and the PRDB would need to continue to support the NSFNET and its regionals well beyond the end of October, a plan was proposed to transition to the RADB to support the NSFNET in its last months.

The new situation presented several problems. First, the tools used to configure the NSFNET/ANSnet routers were based on PRDB attributes, not RIPE-181. Second, the RADB was not yet populated with data. And finally, the PRDB

described AS690 policy with respect to its peer ASs on a per-prefix basis; in the RIPE syntax, the basis for describing routing policy was the Autonomous System where the route originated. With more than 40,000 prefix-based policies for the regionals, the PRDB was used to generate about 100 configuration files of around 250,000 total lines every two weeks, and those policies needed to be re-expressed in a RIPE-compatible format.

Continuing the Policy Routing Database for long-term support of ANSnet was inadvisable. If ANS continued to use the PRDB for AS690 routing after the transition, the PRDB's non-standard format would create a barrier to sharing global routing policies and building tools to aid with global routing. A solution had to be found that would provide stable routing through the transition, and, once the NSFNET was retired, allow the ANS registry to take its place alongside the other registries in the IRR.

To solve the problem, Merit proposed a modification to RIPE-181 -- a temporary attribute that would specify the peer or adjacent AS announcing the route to AS690. The community agreed to Merit's proposal, and the new expression came to be known as the advisory attribute. Merit now needed to quickly modify the PRDB configuration tools so they would generate the new attribute, populate the RADB with the data needed to generate AS690 configuration files, and make sure that the new configurations exactly matched those produced by the PRDB.

By December 1994, all the data in the PRDB had been converted to RIPE-181-style expressions and entered in the RADB. By February, the RADB had been populated with RIPE-181-style Maintainer and AS Objects. The databases were running in parallel, with changes to the PRDB automatically reflected in the RADB. Other organizations whose routing information wasn't related to the NSFNET were also populating the RADB throughout the winter of 1994-95; this was another variable that had to be accommodated as the new database emerged.

Finally came the painstaking task of comparing the config files generated by each database. Merit's Dale Johnson went over the large, quarter-megabyte files line by line, adjusted the configuration tools to compensate for any differences in net lists, and repeated the process over and over until the configs matched perfectly. The RADB finally replaced the PRDB a week after the NSFNET was retired.

Moving the Regionals off NSFNET

NSF and Merit coordinated the process of moving the regionals to new Internet Service Providers, with Merit taking the lead in planning the transition. NSF's new Inter-Regional Connectivity program helped support new attachments not only for NSFNET peer networks -- regionals like SURAnet and NYSERNet that connected directly to the NSFNET backbone-- but also to downstream networks such as NevadaNet and MOREnet. Most of the regionals selected internetMCI or SprintLink as their ISP; CERFnet set up its own ATM connection to each NAP.

In mid-October 1994, NSFNET Program Director Priscilla Huston sent a letter to the regionals asking them to send a transition calendar and engineering overview to Elise Gerich of Merit and to her. Huston also asked the regionals to notify her if they weren't going to make the October 31 deadline for moving off the NSFNET.

As it turned out, none of the networks made it. MOREnet, one of the downstream regionals, missed by only a day; other networks slipped by as much as three or four months. The first NSFNET peer network to make the transition was CA*net, which faced a hard deadline from its link provider for terminating its connection to the NSFNET. The other cutovers were pushed back because of delays in provisioning the ISPs selected by the regionals, and because of reticence on the part of the regionals to move off the NSFNET backbone service.

On one or two occasions, networks that had made the transition had to pull back to full NSFNET connectivity because of deployment problems on the new ISP backbone. In general, though, once the regionals had selected an

ISP and completed all the testing, re-routing, and reconfigurations necessary to make the switch, traffic flowed smoothly over the new infrastructure.

60-Day Notices: No Turning Back

Early in January, when SURAnet notified NSF and Merit that it was ready to move off the backbone, Merit sent ANS the first message to dismantle NSFNET backbone service -- a 60-day termination notice for ENSS 138 in Atlanta. The ENSSs (Exterior Nodal Switching Subsystems) were installed at regional networks attached to the NSFNET, and acted as end nodes for the backbone. This and subsequent termination notices were irrevocable; once sent, there would be no more NSFNET service through that node.

Later in January, NYSERNet and the Cornell Theory Center notified NSF and Merit that they were ready to terminate their NSFNET attachments. The other regionals and supercomputer centers followed suit, one by one, as the April deadline neared. On February 28, Gerich sent Jordan Becker of ANS the formal, 60-day notice for termination of the NSFNET backbone service at 19 locations:

ENSS 128	Palto Alto	April 30, 1995	midnight PST
ENSS 129	Champaign	April 30, 1995	midnight CST
ENSS 130	Argonne	April 30, 1995	midnight CST
ENSS 131	Ann Arbor	April 30, 1995	midnight EST
ENSS 132	Pittsburgh	April 30, 1995	midnight EST
ENSS 133	Ithaca	April 30, 1995	midnight EST
ENSS 134	Cambridge	April 30, 1995	midnight EST
ENSS 135	San Diego	April 30, 1995	midnight PST
ENSS 136	College Park	April 30, 1995	midnight EST
ENSS 137	Princeton	April 30, 1995	midnight EST
ENSS 139	Houston	April 30, 1995	midnight CST
ENSS 140	Lincoln	April 30, 1995	midnight CST
ENSS 141	Boulder	April 30, 1995	midnight MST
ENSS 142	Salt Lake City	April 30, 1995	midnight MST
ENSS 143	Seattle	April 30, 1995	midnight PST
ENSS 144	Moffett Field	April 30, 1995	midnight PST
ENSS 145	College Park	April 30, 1995	midnight EST
ENSS 146	DC	April 30, 1995	midnight EST
ENSS 147	MFS	April 30, 1995	midnight EST

The list included the NSFNET attachments at the NAPs, which were coexistent with ANSnet. ENSS 138 in Atlanta wasn't included, since a termination notice for that node had been issued earlier.

By March, backbone traffic had declined dramatically, but not quite as fast as NSF and Merit had expected. Gerich, concerned that the regionals and ISPs weren't moving fast enough, sent e-mail to the community noting that a significant amount of traffic was still traversing the NSFNET. Merit posted a histogram showing the top 10 originators of traffic into the backbone in February 1995, and reminded networks attached to nodes highlighted on the graph about the April 30 deadline.

Later that month, Merit discontinued the T1 safety net that had backed up the T3 infrastructure since 1992.

Black Friday and the Final Shutdown

By the middle of April, only seven regionals had completely severed their ties to the NSFNET backbone service. Other networks had cut over to a new service provider, but continued to peer with the NSFNET for backup purposes. As the final deadline neared, Merit and the NSFNET Executive Committee became concerned that these redundant connections would make it difficult to identify outstanding reachability issues before the April 30 cutoff.

To spot any pockets of unreachable destinations before it was too late, Merit on behalf of the NSFNET Executive Committee notified the NSFNET community on April 14 that it would terminate peering sessions with all organizations still attached to the NSFNET Backbone service at 9:00 a.m. on Friday, April 21. On April 28, all sessions with the NSFNET service would be permanently terminated; ANS would terminate operation of the NSFNET Backbone service on April 30.

This announcement created quite a stir among the networks attached to the backbone. Several said that they'd lose their Internet connectivity completely if their NSFNET peering was shut down before April 30, and requested that their session stay up. One provider was still relying on his MAE-East NSFNET connection for all his East coast traffic; another requested clemency for a non-production peer router that was proving essential for network diagnostics. A midwest network's installation of a T3 circuit had been delayed; the operators weren't concerned about reachability if their NSFNET peering was shut down, but about capacity—a large volume of traffic was still traversing the NSFNET, and cutting back to a T1 would lead to unacceptable response times. Merit made separate arrangements to accommodate each network, but held to the new deadline.

As it turned out, the test shutdown had to be postponed. In the early morning hours of April 21, Merit notified the community that it would have to delay the regular Friday backbone configuration run. The volume of routing configuration changes had increased so dramatically as networks switched to new providers that some of the files grew large enough to truncate during production, and produced corrupt configuration files. Merit wasn't confident that it would be able to produce complete and correct configurations in time for the normal 8:00 a.m. configuration window. Additional file space had to be allocated before the configs could be run, and Merit needed to work with several ASs to reduce the number of net lists in the config file. This meant postponing the Friday shutdown until Saturday, and delaying the NSFNET discontinuation until Tuesday, April 25. The test shutdown had indeed pinpointed at least one problem as a result of delayed transitions: the processing of several thousand simultaneous changes to router configurations was more than the PRDB could handle.

On Monday, one network jumped the gun, and surprised Merit and ANS by turning off its ENSS. The ANS staff noted that no harm had been done, but reminded the sysadmin that the plan was to manually turn off peering on the 25th, and shut off the ENSSs on the 30th. IBM was to physically remove the routers beginning in May.

On April 25, the peering sessions on 15 ENSSs were commented out of the configuration files and the NSFNET Backbone Service was, for all intents and purposes, terminated. The next Sunday evening at midnight, a dozen or so staff from Merit and ANS gathered in the University of Michigan NOC to turn off the ENSSs, one by one, at midnight in each respective time zone. One or two regional operations centers called the ANS NOC about unreachable ENSSs, but "mostly the NSFNET went away silently," as one ANS engineer remarked, "or rather, with only the sound of drives and fans spinning down in distant machine rooms."

On May 8, with Merit confident that the RADB was producing consistent configuration files for the ANSnet and ANS ready to take over configuration generation for AS690, the PRDB made a graceful exit. The new architecture was in place: internetMCI and SprintLink had absorbed the NSFNET regionals as their customers; the RADB and the databases maintained by the RIPE NCC, internetMCI, CA*net, and ANSnet had replaced the PRDB as a means of describing routing policy.

Farewell NSFNET! And congratulations to the hundreds of people who helped make the backbone such a great success.

Footnotes

1. Frazer, Karen D. "NSFNET: A Partnership for High-Speed Networking.", Merit Network, Inc., 1995.
2. Ibid.
3. Ibid.
4. Ibid.
5. Bates, T., Gerich, E. Joncheray, L., Jouanigot, J-M., Karrenberg, D. Terpstra, M. and Yu, J. "Representation of IP Routing Policies in a Routing Registry (ripe-81++)", October 1994. Subsequently published as RFC 1786, March 1995.

About the Authors

Susan Harris coordinated NANOG meetings, chaired the NANOG Program Committee, and moderated the NANOG email list. A Senior Science Writer at Merit, Susan provided user support for various operations and research projects and for Merit's regional backbone, MichNet. She was the author of three RFCs. Prior to joining Merit in 1992, Susan earned a Ph.D. in Near Eastern Studies from the University of Michigan.

Elise Gerich was a product manager at Juniper Networks. Prior to joining Juniper in 2001, she held various management and technical positions at Urban Media and Excite@Home, where she helped roll out the company's first broadband network. Prior to Excite@Home, Elise was Associate Director for National Networking at Merit Network, Inc. In this role, she worked closely with the U.S. regional networks to ensure a smooth transition from the NSFNET Backbone Service, served as Co-PI of the Routing Arbiter Project, and founded NANOG with colleague Mark Knopper. She is a former member of the Internet Architecture Board and a longtime participant in the IETF, NANOG, and RIPE.