# Link Letter

The Merit Computer Network

Vol. 1 No. 2 — 15 April 1988


## *The Link Letter* starts new electronic edition to meet subscription demands

This is the second issue of *The Link Letter*. For those of you new to the publication, *The Link Letter* is produced by Merit and focuses on issues of interest to the mid-level networks and supercomputer centers directly connected to the National Science Foundation network backbone. As Merit works on network re-engineering and puts into place a new management system for the NSFNET backbone, *The Link Letter* will provide timely information on technical developments and scheduling information. Future issues will also report on operations and ongoing network research related to NSFNET development.

The response to our first issue has been overwhelming and we appreciate the many positive comments we received from readers. Because we have had more requests for the newsletter than we ever expected, we are implementing a new circulation plan to meet the demand. Starting with this issue, most subscribers will be getting the text of *The Link Letter* electronically. Multiple hard copies will be sent to each of the thirteen remote sites and to a few other locations directly affected by the NSFNET backbone re-engineering. Anyone interested in receiving the electronic version of *The Link Letter* should send a message to:

NSFNET-Linkletter-Request@merit.edu

For general information on the NSFNET project, you can contact:

NSFNET-Info@merit.edu

or call (313) 764-9430. We welcome suggestions about articles you would like to see in future editions.

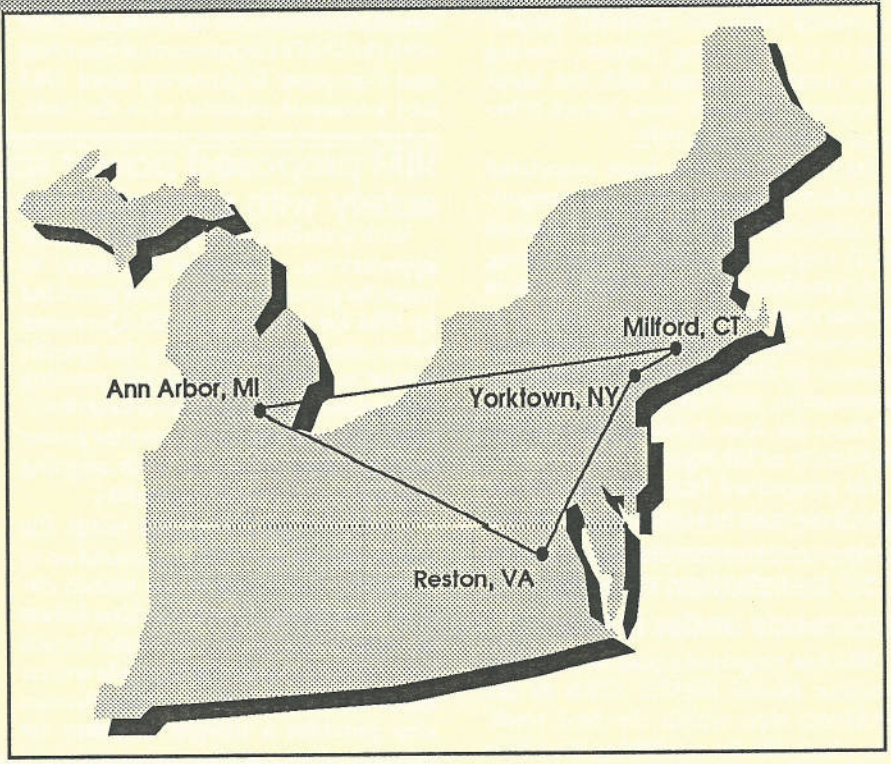


## TCP/IP testing under way on four node research network

The four node experimental network that we are using as an engineering testbed for the NSFNET backbone network has full physical connectivity. The packet switches (NSSs) at Milford (IBM), Ann Arbor (Merit), and Reston (MCI) are fully operational. The NSS at IBM Yorktown will be connected in the very near future to complete the four node logical topology. Currently, a prerelease of the NSS software is being tested over this configuration. The test network is operating at one-third T1 speed (448 kilobits per second). Almaden cards are being used for the synchronous RS422 interface. Plans are under way to switch to cards capable of higher speeds in the future.

In addition, we are carrying out preliminary tests to determine approximate performance figures. The effect of the IDNX as a circuit-switching component is also being investigated.

We will be testing the first version of the routing software on the experimental network shortly. This is the IBM-NSFNET implementation of the ANSI SPF-based routing protocol. This will be used for backbone routing in the operational NSFNET.

*–Bilal Chinoy*

Introducing Eric Aupperle

# Merit director brings long-term experience to NSFNET backbone development project

As Director of the Merit Computer Network, Eric M. Aupperle also serves as the Project Director for the NSFNET backbone re-engineering program. Eric directly supervises all Merit's activities associated with the NSFNET project. Eric is a networking pioneer, having been deeply involved with the Merit Computer Network even before it became operational in 1972.

Eric Aupperle has been associated with electrical engineering and computers since the early 1960s. He received B.S.E. degrees in electrical engineering and in mathematics, an M.S.E. degree in nuclear engineering, and the Instm. E. (a professional degree in engineering granted by the Program in Computer, Information, and Control Engineering) all from the Ann Arbor campus of the University of Michigan.

He progressed from an assistant research engineer to a research engineer at Cooley Electronics Laboratory, the University of Michigan, where he worked with fuse countermeasures, applications of solid-state devices, and pseudo-random communications system countermeasures. He has been a lecturer with the U-M's Department of Electrical and Computer Engineering since 1963 and a research scientist with the Institute for Science and Technology since 1969 at the University of Michigan.

In addition to serving as Director of the Merit Computer Network and of the National Science Foundation's NSFNET backbone network, Eric is also Associate Director for Network Services at the University of Michigan Computing Center.

## CSU installation at sites proceeds under IBM plan

IBM has proposed a plan to install the Verilink Model 551VST CSUs at the backbone sites within the next week. The plan calls for modifying the CSUs and putting them in place until UL-replacements are available.

IBM is coordinating the efforts of Verilink, MCI, and Merit toward this end. The first step of the plan proposes that Verilink provide safety shrouds for the CSUs that are currently at the sites. Merit received the shrouds on April 11 and will express mail them to sites by April 15. Installation of the shrouds resolves the safety concerns for the short term. MCI will modify and install the CSUs and run end-to-end testing of the physical network starting April 18.

We anticipate completion of the end-to-end testing within two weeks. As UL-approved replacements are available, IBM will replace the field-updated Verilink CSUs. We are hopeful that this plan will provide equipment and service to each site with as little disruption as possible.

## IBM proposed power specs for NSSs promote safety with potential for future expansion

With installation of the NSSs rapidly approaching, most sites are ready to meet the power specifications provided by IBM for running an NSS. Questions have been raised regarding the specific power requirements and the seeming redundancy in the number of branch circuits. The need for the particular power configuration relates to both ongoing operation and future expansion.

Under the present power setup, the remainder of an NSS will continue to function even if one of the branch circuits is lost because of a machine failure or power outage. The multiple branch circuits help ensure that the node will remain intact. The power specification also provides a margin of safety by limiting the effects of power variation on individual pieces of equipment.

In looking toward the future, these power requirements allow for network growth and development. The present specification provides sufficient power for NSS expansion at each site. This is a proactive system aimed at installing sufficient power now rather than reacting under stress in the future when demand exceeds the power allocation. In terms of cost, putting the additional capabilities in initially is less expensive than the cost of future expansion. If you have any questions about power requirements, message:

NSFNET-Info@merit.edu

or call (313) 764-9430 for information.

## NSFNET News Briefs . . .

**Delivery of IDNXs.** IBM has begun shipment of IDNXs to the backbone sites. The IDNX 20 has been tested by IBM safety and granted a corporate safety deviation. This means that the IDNX 20 presents no safety hazard when installed in controlled-access machine-room environments. The National Service Division of IBM will install an IDNX at each site and will provide service to it. Merit will be contacting each site to verify the arrival of its IDNX and to schedule a date for its installation.

**NOC machine in place.** The two IBM 4381 mainframes housed in the Ann Arbor Network Operations Center are now in their permanent positions in the remodeled Computing Center Building on the University of Michigan campus. The 4381s will be used for network operations control and for information services on the NSFNET backbone. The operational NSS for the Merit site has been assembled and moved to its permanent location in the NOC.

**Information Services Update.** Richard Schmalgemeier recently joined our staff as the Online Systems Coordinator for Merit/NSFNET. Rick is very experienced in implementing and managing mainframe information services. Rick, with Jim Sweeton, manager of Merit/NSFNET Information Services, attended the SPIRES consortium in Palo Alto in March, and visited SRI-NIC.

# Link Letter

The Merit Computer Network

Vol. 1 No. 3 30 April 1988


# Physical connections for NSFNET backbone ready for operation

MCI has completed installation and end-to-end testing on all the T1 lines which form the physical NSFNET backbone. From one to four circuits connect each of the thirteen sites to MCI Points of Presence (POPs). In the center of this configuration are two interconnected rings using a combination of high quality digital radio and optical fiber for the transport of digital data between POPs. Located along the rings are Digital Cross Connects (DXCs) which enable circuit switching to reroute data.

No single line failure or outage will disrupt service on the rest of the network (although individual sites that are singly connected could be without service if their circuit to the nearest POP failed). This configuration not only provides redundancy to assure network reliability but also accommodates additional nodes in the future.

The software-controllable DXCs will play a major role in restructuring transmissions to provide high reliability. The DXCs allow real time control over bandwidth and configuration of the circuits. If an outage is detected on a line, the traffic can be rerouted automatically through unaffected lines in the network. In the early phases of NSFNET re-engineering, software-controlled DXC circuit switching will not be fully implemented. Limited monitoring and restructuring of transmissions using the DXC technology will be done by MCI. Later phases will turn this circuit-switching capability over to the Merit/NSFNET Network Operations Center (NOC) for rapid problem resolution. The circuit-switching system planned by Merit and MCI is flexible enough to incorporate new technologies as they become available.

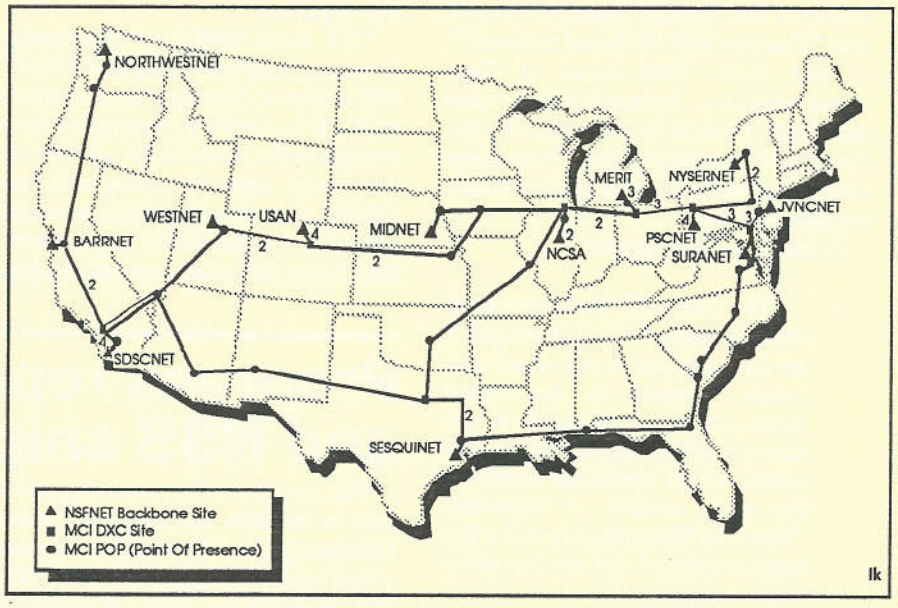


Better understanding of the role circuit switching plays in the re-engineered backbone is gained by looking at the network hierarchy as planned by Merit. At one level is the physical topology consisting of the interconnected rings and backbone spurs providing connectivity to the thirteen backbone nodes. At the next layer is the interconnectivity of the circuit switches provided by the DXCs in MCI's Digital Data Network (DDN), allowing real time reconfiguration of the physical connections. Above these are the logical interconnections of the packet switches composed of the links between Packet Switching Processors (PSPs) in the Nodal Switching Subsystems (NSSs).

Although dynamic circuit-switching capability using DXC technology is central to the future of the NSFNET backbone, the Phase I network installation scheduled for July relies more heavily on IDNX technology for determining bandwidth and connection restructuring. These programmable multiplexors will be controlled from the NOC to adjust for line failures or localized traffic congestion. Early stages of testing and statistics collection will be used to implement faster and more efficient methods of restructuring in future operations. A critical element is to develop these methods without impeding normal network operations.

*—Ellen Hoffman*

*(Editor's note: The next issue will have an article covering the role of IDNX technology in Phase I backbone implementation.)*

## Newsworthy NSFNET responds. . .

# "Kludgey" depends on one's field of vision

Articles featuring NSFNET appeared in recent editions of *Infoworld* and *Data Communications*. With the re-engineering project heating up as the July cutover approaches, the national press is paying increased attention to network developments.

Unfortunately, a number of inaccurate statements appeared in the articles "The Largest Computer Network" by Sharon Fisher in the April 25 issue of *Infoworld* and "Kludgey NSF RT-based network?" in the April 1988 issue of *Data Communications*. *Infoworld* did not contact Merit prior to publication. We believe that *Infoworld* and *Data Communications* neglected several important details in these articles.

In response to *Infoworld*:

- The NSFNET backbone links thirteen sites: six supercomputer centers and seven mid-level academic networks in the United States. Canada's Netnorth is not part of the NSFNET backbone. (An article which appeared earlier in the same edition of *Infoworld*, "National Research Net to be Enhanced in July," did get this information correct.)
- The NSFNET award for re-engineering and managing the backbone went to MERIT, not the University of Michigan, IBM, or MCI. MERIT, Inc. is the consortium of eight universities in Michigan which manages the Merit Computer Network. Merit has established joint study agreements with IBM and MCI on network research and development. As part of this effort, MCI is leasing lines for the backbone and providing expertise in circuit-switched technology. IBM is providing hardware and software for NSFNET packet-switching nodes and network management.

The article which appeared in *Data Communications* has resulted in a number of concerned messages from individuals associated with NSFNET. IBM's Jack Drescher wrote the following response:

"The April 1988 *Data Communications* magazine infers that the NSFNET NSS nodes are 'kludgey' combinations of PC/RTs. I believe the article fails to recognize the modularity, on-line redundancy, network management capability and orderly growth potential provided by the initial implementation. It also fails to acknowledge a production network requirement met by virtually all standard, proven hardware in a timeframe of about one year from the publishing of the requirement.

"We wrote to you in *Link Letter* (Volume 1, Number 2) citing that extra reliability and orderly growth potential were heavily considered in IBM's power specifications for the NSS nodes. I think the basic design, architecture, and implementation lend themselves to technology and packaging enhancements as they naturally develop.

"The article also states that *Data Communications*' calls were not returned by me in time for publication. I have no record of any calls from *Data Communications* to me on this subject."

*—Ellen Hoffman*

# Progress report on IDNX and NSS shipping

The first five NSSs were shipped to sites in Boulder, San Diego, Princeton, Champaign, and Ithaca. MCI also completed testing of four circuits, including Princeton-Ithaca; San Diego-Champaign; Palo Alto-San Diego; and Salt Lake City-Boulder.

All sites should receive completely configured NSSs by mid-May. The frames for the components are already at each of the sites.

IBM spearheaded a plan to modify and install the Channel Service Units (CSUs) to conform to safety standards. Because of the cooperative efforts of IBM, MCI, and Verilink, the current model of the CSU was modified and end-to-end testing of the T1 links is now in progress.

Pittsburgh, Lincoln, Ithaca, and Ann Arbor have received their IDNXs and will be contacted shortly by IBM to schedule the installation and assurance testing of this equipment. Shipments of the other IDNXs are scheduled to begin on April 29th. The target is to have all IDNXs installed by the middle of May.

Thanks to everyone for fulfilling our requests in a timely fashion. With the completion of the installation process, we are about to embark on the next phase of the project: testing of and migration to the new NSFNET backbone network.

*—Elise Gerich*

*The Link Letter* is published by Merit/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of *The Link Letter* should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions about articles you would like to see in future editions. Staff: Elise Gerich, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Jim Sweeton.

# Link Letter

The Merit Computer Network

Vol. 1 No. 4 13 May 1988


# IDNX technology plays crucial role

*(Editor's note: This is the second article in a series on dynamic reconfiguration in the NSFNET backbone. As noted last week, Merit is working with MCI to incorporate leading-edge technology into the physical network. This article highlights IDNX technology which will play a critical role in the first phase operational backbone.)*

The circuit switching equipment used to implement Phase I in the NSFNET backbone project is an intelligent T1 communications processor called an IDNX (Integrated Digital Network Exchange). Each NSFNET backbone site will have an IDNX. The models vary among IDNX 20s , 40s, and 70s, depending on the number of physical connections a particular site has to its Nodal Switching Subsystem from the backbone.

The IDNX is a software-controlled circuit switch provided by IBM. It was initially designed to maximize T1 trunk transmission efficiency and manage hybrid networks combining private and public data networks. In this mixed environment, the IDNX capability to restructure transmissions without powering down the system is a major advantage to network operations.

For the new NSFNET backbone, the physical topology comprises thirteen IDNXs interconnected by T1 trunk lines. The IDNXs provide three capabilities to the network: subrate T1 multiplexing, dynamic alternate routing, and dynamic bandwidth allocation.

Subrate T1 multiplexing allows the attached Packet Switching Processors (PSPs) in the NSS to operate at any data rate up to full T1 (1.544 mbps) rates. By multiplexing sub-T1 data rates on a single T1 channel, the number of local links (PSP connections) is not limited by the number of physical circuits to the NSS. In other words, one physical circuit is divided as though it were multiple connections. This is shown in the diagram. In the link from IDNX B to IDNX C, the circuit is "divided" in half.

A second feature is the ability of IDNXs to reconfigure connections in the event of link failure. If the A-B link failed, transmissions from A could still get to B by restructuring the circuit flow to go A to D to B. Limited failure of an IDNX will affect the network only briefly, as an alternate route will be chosen from among the remaining IDNXs.

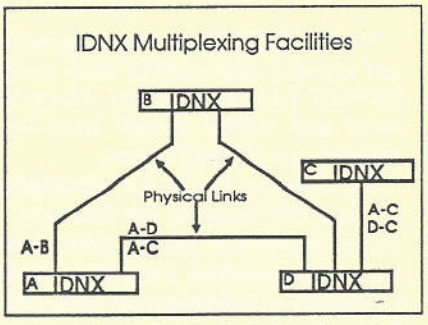


As the system demands change, the NOC can use the IDNX to allocate bandwidth on the T1 trunk lines. This dynamic bandwidth allocation permits an increase or decrease to the amount of bandwidth allocated to any PSP in response to traffic patterns. For example, in the B-C link, the A-D could use 3/4 of the total bandwidth while A-C may need only 1/4.

The use of IDNX technology for network management is currently being evaluated and refined on the four-node research network. With the first IDNXs now installed on the production NSFNET backbone, further tests can be performed before the July cutover.▼

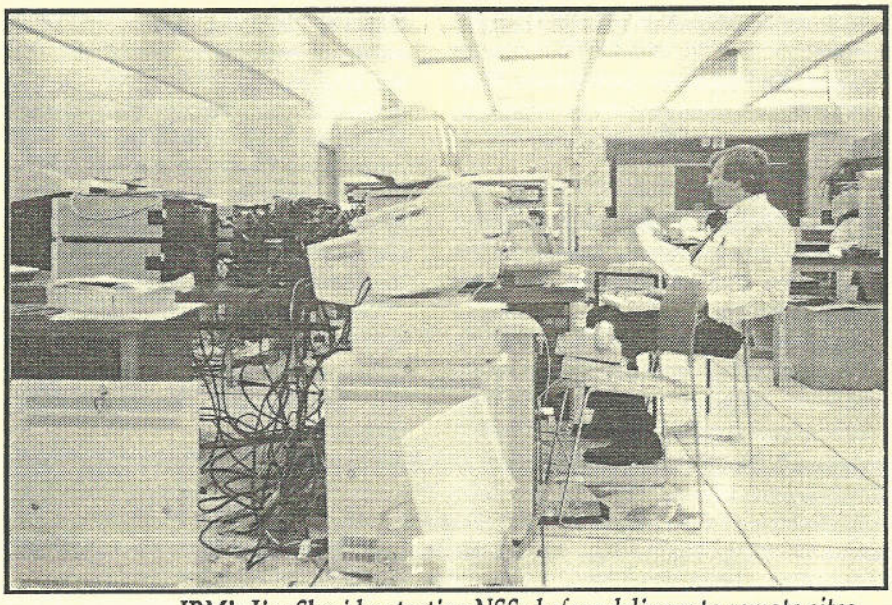

*-IBM's Jim Sheridan testing NSSs before delivery to remote sites.*

# Meet NSFNET's Principal Investigator

Hans-Werner Braun is Principal Investigator for Management and Operation of the NSFNET Backbone Network. He is also heading the Internet Engineering group of the Merit Computer Network and is Principal Investigator on an additional NSF award to study routing issues for the NSFNET.

Hans-Werner was born in West Germany. After receiving his engineering degree, he worked for five years in data communications at the Regional Computing Center of the University of Cologne. During this period, he also did consulting for a firm developing and installing a networked reservation system of a world-wide hotel chain. Hans-Werner joined the University of Michigan's Merit Computer/UMnet group in April 1983.

Even before the project began at the University of Michigan, Hans-Werner

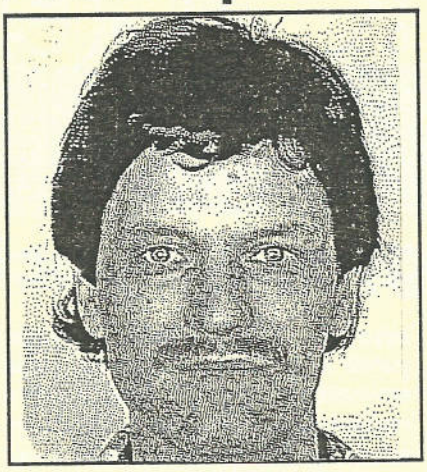

participated in the design, implementation, and operation of the NSFNET backbone as well as of various mid-level networks to be attached to the backbone.

Active in a number of professional societies, Hans-Werner is a member of ACM Special Interest Group on Communications, and of the IEEE Communications and Computer Societies. He formerly chaired the Technical Committee of the NSF/DNCRI Network Program Advisory Group. Hans-Werner is also a member of the Network Technical Advisory Committee of the John von Neumann National Supercomputer Center. He participates in a number of other committees and Internet Task Forces, as well as the Steering Group of the Internet Engineering Task Force.

Within the current project, Hans-Werner chairs the Joint Network Technical Committee and, together with Joe Choy of the National Center for Atmospheric Research, is responsible for the NSFNET Network Working Group. This group exists on behalf of both FARNET, which Joe represents, and the NSFNET Backbone Project. ▼

# Team building progresses at Merit/NSFNET

Merit/NSFNET has been busy hiring and training new staff to work on the NSFNET backbone project. Many talented and skilled people have been added to the NSFNET team.

Jim Sweeton heads the Information Services (IS) team as the Manager of Merit/NSFNET IS. Jim has been associated with academic computing at the University of Michigan for five years. Before that, he was a systems manager at the University of Michigan Hospitals. His particular interests are in large information systems and in the effective distribution of information in a networked environment. Jim is active in the local homeless shelter and food program activities in Ann Arbor, and participates in the yearly "Pedal Across Lower Michigan" bicycle trip.

Also in the IS group is Elise Gerich, who is a technical site liaison providing consulting for staff at the thirteen NSFNET backbone sites. Elise previously worked at Network Technologies International, Inc. (NETI), a small software development firm marketing UNIX and XENIX based products. Elise worked as the Beta Site Manager with such organizations as Mitsubishi Electronic, TRW, FDA, and AT&T. If you are one of the backbone site contacts, Elise has probably already been in communication with you either electronically or by telephone.

Laura Kelleher recently joined the Merit/NSFNET IS team as a technical writer. Laura previously worked for the Residence Halls Computer Program (RESCOMP) at the University of Michigan where she was Senior Trainer. Laura will be working on various Information Services projects, such as producing articles and graphics for the Link Letter, writing technical documentation, and creating public relations materials.

Susan Hares has accepted the position as the Routing Coordinator for NSFNET within the Internet Engineering group. She will start at Merit on May 23, but is already reachable via electronic mail. Susan has more than ten years of technical work experience and is currently working as a Senior Software Engineer in the ISO/MAP Network Management Department of Allen-Bradley's Communication Division. Prior to working at Allen-Bradley, Susan was a Senior Systems Programmer in the Network Development Software Department of the Autonet Division of ADP.

We are all very excited about these new members of the NSFNET/Merit team and we look forward to their further participation on the backbone project. ▼

The Link Letter is published by Merit/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of The Link Letter should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions about articles you would like to see in future editions. IS Staff: Elise Gerich, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Jim Sweeton.

# Link Letter

The Merit Computer Network

Vol. 1 No. 5 27 May 1988


# Seven sites online and operational

Seven Nodal Switching Subsystems (NSSs) are installed and online on the backbone as of May 31. These NSFNET sites are Ann Arbor, Boulder, Ithaca, Lincoln, Palo Alto, Princeton, and San Diego. The other six are on the test network at Yorktown Heights (IBM), Reston (MCI), Milford (IBM), and Ann Arbor (Merit). The two network partitions are linked in Ann Arbor.

Over the next week, the remainder of the backbone sites are scheduled to come online. By June 10, all 19 NSSs should be installed and exchanging IS-IS routing packets. The project remains on schedule.

**The Software Connection**

Central to full connectivity for the NSFNET backbone is the software developed for the NSSs. This software not only must provide routing from NSS to NSS, but outward to the each of the primary "campus" level or local networks that are attached to the mid-level networks. The Routing Control Processor (RCP) within the NSS plays a primary role in mediating the routing information needed by the local networks and that needed by the NSFNET core. The RCP can collect routing information from a regional gateway attached to the NSS and broadcast that data to other RCPs on the backbone. Work presently underway and continuing over the next month is aimed at refining the connectivity between the backbone and regional networks.

For routing within the NSFNET core, the NSS code incorporates the IBM adapted ANSI Intermediate System to Intermediate System (IS-IS) Intra-Domain Routing Exchange Protocol, which uses a Shortest Path First (SPF) algorithm for internal routing. These IS-IS packets are now being exchanged between the operational NSSs.

With the NSSs online, testing is proceeding on improving connectivity to the local networks. Such connectivity has been established between the new backbone and several machines emulating gateways to the local networks.

Currently, gateways at Cornell University (Unix), SDSC (VMS), NCAR (Unix), and JVNC (Cisco) are running the Gateway Daemon (gated) developed at Cornell University. In addition, initial tests were done using Proteon gateways.

The next step in phasing in the "new" network is to improve the routing interface between the RCPs and other local gateways via the External Gateway Protocol (EGP). This testing is being done as more mid-level networks are in the process of adding Ethernet connections between the NSSs and their local networks. Plans are to begin testing "live" traffic in early June as the local networks come online.

**Network Management**

In addition to NSS development, the network management system for the new backbone is being readied for operation. Netview and Netview/PC have been installed and configured at Merit's Network Operations Center (NOC) in Ann Arbor. This software plays a major role in alert monitoring and display for the new backbone. Testing is proceeding on a Simple Gateway Monitoring Protocol (SGMP) implementation which will be used to monitor network traffic, routing tables, and link status on the backbone. SGMP testing is being done on the experimental network by researchers at Yorktown Heights with the assistance of Jeff Case and Ken Key of the University of Tennessee.▼

*This article includes information contributed by Jordan Becker, IBM-Yorktown.*

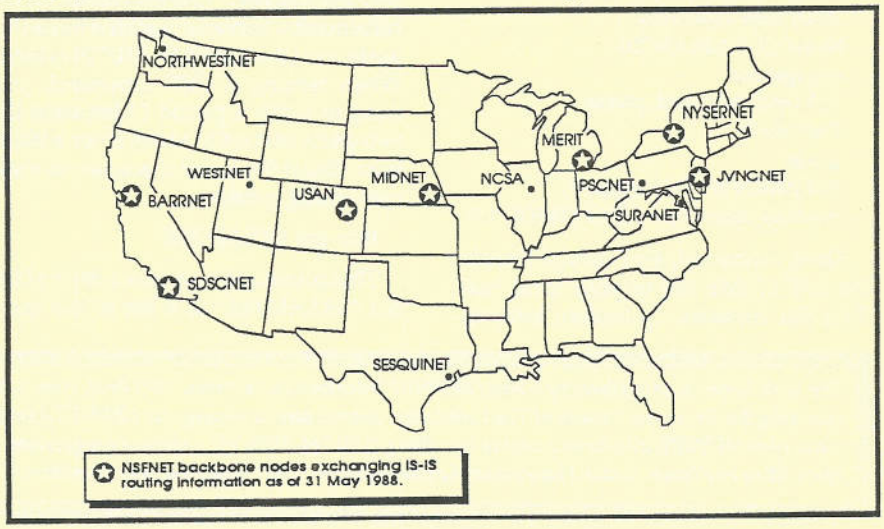

# FTP now on Information Services machine

File Transfer Protocol (FTP) is now operational on the Merit/NSFNET Information Services machine. In general, FTP allows users to communicate with a remote host and to transfer both text and binary files between a local machine and a remote machine. FTP service on the IS mainframe allows you to retrieve files related to the NSFNET project.

The availability of files for FTP is one of Merit's efforts to meet the informational needs of the NSFNET user community. As the IS effort develops, a number of documents will be available, including the *Link Letter*, Merit User Memos, NSFNET research reports, and information on network operations, management, and performance.

To transfer files using FTP, first invoke the FTP program on your local computer. To access the remote machine, open a connection to the remote host by machine name or Internet address. In this case, the remote machine is 35.1.1.48 (the Internet address of the Information Services Machine) or NIS.NSF.NET (its name).

Once you have a connection to the IS machine you will need to log on. IS provides an "anonymous login" capability. The example FTP session below indicates the sequence of commands used to transfer files:

```
ftp> open NIS.NSF.NET
220-FTPSERVE at NIS.NSF.NET,
 220 Connection will close if idle for
 more than 5 minutes.
Name (35.1.1.48:ABCD)
anonymous
 331 Send password, please
 Password:
 guest
 230 ANONYMOU logged in;
 working directory = ANONYMOU 191
```

Once connected, you should transfer the $READ.ME file to your local host. This file contains up-to-date information about the various directories on the Information Services machine.

The following directories are currently available on the IS machine: NSFNET, LINKLTTR, and MERIT. The NSFNET directory contains research reports pertaining to NSFNET activities. Copies of the electronic version of the *Link Letter*, are in the LINKLTTR directory. The MERIT directory contains the Merit User's Memos and related information files.

To view the contents of a directory, issue the DIR command. To change directories, use the CD command.

```
ftp> dir
200 Port request OK
125 List started OK
$INDEX LINKLTTR 26  ANY191
$INDEX MERIT   173  ANY191
$INDEX NSFNET   26  ANY191
$READ  ME      101  ANY191
                ...
ftp> cd nsfnet
250-Workingdirectory is NSFNET
ftp> dir
200 Port request OK
125 List started OK
$INDEX  NSFNET   25   NSF191
$READ   ME       65   NSF191
BRAUN   ROUTING 463   NSF191
REKHTER EGP     265   NSF191
```

If you find an interesting file, you can transfer it to your workstation or local host to read it. The GET command initiates the transfer. The file names on the Information Services machine have two parts, as in the "$READ ME" file above. When issuing a GET command, you need to insert a period (.) between the two parts of the file name, since a blank is an illegal file name character on most computing systems.

```
ftp> get $READ.ME
```

This command transfers a copy of the file "$READ.ME" to a file of the same name on the user's local disk space. For more information about FTP, enter HELP at any prompt.▼

# An open invitation

Merit/NSFNET would like to extend an invitation to the individuals responsible for routing concerns for their regional networks to come to Ann Arbor to meet with the Merit/NSFNET team. Meeting here will provide an opportunity for members of the NSFNET backbone sites to exchange information with various members of the NSFNET backbone and management operations center here in Ann Arbor.

Possible areas of discussion for these visits include: network monitoring for troubleshooting problems, routing issues concerning different configurations of the architecture, information services needs, and escalation procedures. You can also take a tour of the NOC.

Four sites have already taken advantage of this opportunity. Dale Finkelson (University of Nebraska-Lincoln) and Ken Lowe (University of Washington) visited Ann Arbor to help with the assembly of the NSSs for NorthWestNet and Midnet, as well as to discuss routing for those sites. Guy Almes discussed Sesquinet's regional topology and network monitoring needs. Don Morris (NCAR) was also here to talk about routing and other issues. BARRnet's David Wasley and Abraham Bleiberg are scheduled to visit in early June.

If you'd like more information about such a visit, call us at 313-764-9430 or at NSFNET-info@merit.edu. ▼

The Link Letter is published by Merit/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of The Link Letter should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions and contributions. IS Staff: Elise Gerich, William Hart, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Andrew Simms, Jim Sweeton.

# Link Letter

The Merit Computer Network

Vol. 1 No. 6 10 June 1988


# GRASPing for information

As many of you know there comes a time during the intense final stages of putting together a large grant proposal when the weary participants move to heights of zaniness otherwise rarely encountered. It was this aura that inspired our name for the vehicle that will ultimately handle all aspects of electronic distribution of information for the Merit/NSFNET team and the NSFNET community — GRASP.

GRASP is a combination of a front-end system (GRANDiose Distributed Application System or GRAND) and a powerful database management system back-end, SPIRES (Stanford Public Information Retrieval System). Together they will provide both standard mechanisms as well as fully-sophisticated functionality to serve our users.

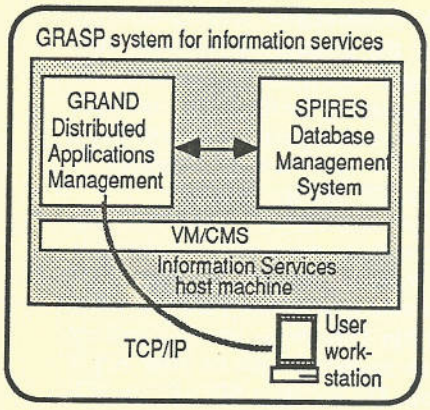


GRAND is a sophisticated server which has been in operation over the vast IBM Corporate Network, VNET (2,000 major computers in some 215 cities throughout the world), for several years. Some of the GRANDiose design considerations exemplify GRAND:

•GRAND supports a variety of communications applications in a "normal electronic environment." Since no single user interface is satisfactory for all users, the individualization of the interface is reasonably easy in the distribution environment

•Authentication of users in a distributed environment is a highly desirable attribute of systems that exchange information. This is particularly critical when some or all of the information is tagged with the poster's name (mail, computer conferencing, for instance). For those applications where such authentication is desirable or mandatory, a GRID (GRAND ID) and associated password are required before the application becomes available to the user.

•Neither central- nor local-server independence is desirable, but complete "cooperation" and concurrence between distributed and central servers is essential. GRAND is quite intelligent in this regard, and there is considerable experience in server-to-server interactions. This distribution of resources among peered servers, with central GRAND management, will be efficient in our backbone environment. We are hoping that some of the backbone sites will be interested in running local GRASP servers.

The GRAND system is based on components called "styles." Styles are both transaction-types as well as the associated software, and GRAND has the ability to define other styles as needed. Fundamental styles include:

•Notification Style (NS) is the transaction that informs the user when a change in an application or entity (file, conference, database) occurs. A common use of NS is for the GRAND computer conferencing application, where the GRAND user (or another GRAND server) wants to know if some other user (perhaps via remote GRAND server) has added comments or a discussion topic.

•Conference Style (CS) is a higher level function than NS, and its job is to manage a particular computer conference model. In the GRAND scheme of things, the concept of a "computer conference" is not a fixed entity, but rather a template for assisting in the creation of an interactive (and perhaps distributed) application that fits the particular needs of the group wishing to communicate.

•User Interface Style (UIS): in a full implementation, UISs are combined to form User Interface Programs. The UISs run in that "user's normal computing environment" that might be anything from VM/CMS to a small microcomputer workstation. This is the front-end of the front-end GRAND system, and coordinates the transactions between the user and the local server, as a two-way stream of styles.

So, for us at Merit/NSFNET Information Systems, GRAND represents a sophisticated set of high- and low-level tools to make our job considerably easier, and to give us a basis for standardizing information-exchange activities in our NSFNET backbone environment. We will be sharing our trials, tribulations, successes, and failures with the GRAND and SPIRES communities, and we look forward to your comments, suggestions, complaints, compliments, and support in our online systems endeavors!

*Note: This is the first of a three-part series on the GRASP System.*

# Four more join NSFNET Information Services

The Merit/NSFNET staff continues to grow as work on the backbone progresses. Information Services has four additions to its team.

Bill Hart and Andrew Simms, both University of Michigan students, are new computing consultants working for Merit full time this summer and part time during the next academic year. Bill is a senior majoring in Honors Mathematics. For the past two summers, he worked as an engineering aid at Kelsey Hayes Corporation in Brighton, Michigan. Andrew Simms is a senior majoring in Computer Science. His past experience includes three years as a part time employee in the Computing Center Communications group. Bill and Andrew are assisting Rick Schmalgemeier with writing on-line help facilities for NSFNET, working as UMnet postmasters, and creating database applications.

Janise Honeyman recently became a permanent member of the Information Services team as a technical writer. Janise comes to Merit/NSFNET with five years of experience in software documentation. She has worked as writing consultant at University of Michigan Business School, and has taught technical writing, creative writing, and English composition. After completing an M.A. in English, she continued with doctoral studies.

The newly hired coordinator of the Information Services Communications group is Ellen Hoffman. As a staff member of the Merit Technical Support Group, Ellen has been assisting with the NSFNET project since last November but is now permanently assigned to the Merit/NSFNET team. She came to Merit from the University of Michigan Computing Center, where she served as a user consultant and educational coordinator. Ellen has M.A.'s in journalism and anthropology, and has completed coursework towards a doctorate in anthropology with a specialization in archaeology and ethnobotany.

# Link Letters from the backbone

**To: nsfnet-info, nsfnet-info@MERIT.EDU**
**From: steve@note.nsf.gov**
**Subject: I tried to ftp...**

$READ.ME from nis.nsf.net per Link Letter V.1 #5, and got

*$READ.ME: Permission denied*

(Issuing the dir command didn't show anything resembling that filename; is it really there?) Am I missing something? Also, how can I cd back to the top level? On a UNIX system,

*cd ..*

would do the trick - but that doesn't work here.

-s

Steve, thanks for your excellent debugging of our *Link Letter* FTP article. In the sixth paragraph, we say 'you should transfer the $READ.ME file...' but at the time you tried to FTP there were no $READ.ME files on any of the directories. This oversight has been corrected and the $READ.ME file now exists on all of the directories.

The file system under VM is not hierarchical, so the concept of a root or top level directory is not applicable. To get back to the "home" directory for anonymous FTP, a

*cd anonymous*

command will do what you want. Otherwise things worked as advertised in the article.

Since FTP on the IS machine is running under VM on an IBM 4381 mainframe, a couple of other problems have arisen for some users. In the VM/CMS context, filenames starting with a '$' are legal (and force a file to show up at the top of a DIR listing where alpha collation is the default). On some systems, a '$' at the start of a filename may not be legal, or may have a different interpretation. For instance, on a UNIX™ system the leading '$' is interpreted to mean that the following string is the name of a shell variable. Since various systems have different file naming conventions you may need to use the following form of the FTP "get" command to assign a name compatible with your local system when transferring a file:

*get <remote.file> <localfile>*

where "<remote.file>" is the name of the file on the Information Services machine, and "<localfile>" is the name you are transferring the file to on your host. To avoid extra typing, we have put up a READ.ME file that is identical to $READ.ME. We aim to please.

People who FTP to/from IBM mainframe systems need to remember that these machine are EBCDIC-based, and you should <u>not</u> do binary file transfer of a text file. Otherwise, you will end up with an unreadable file on your ASCII system. Do not type <u>BINARY</u> in FTP unless you are transferring executables or other non-text files.

The Link Letter is published by Merit/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of The Link Letter should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions and contributions. IS Staff: Elise Gerich, William Hart, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Andrew Simms, Jim Sweeton.

# Four more join NSFNET Information Services

The Merit/NSFNET staff continues to grow as work on the backbone progresses. Information Services has four additions to its team.

Bill Hart and Andrew Simms, both University of Michigan students, are new computing consultants working for Merit full time this summer and part time during the next academic year. Bill is a senior majoring in Honors Mathematics. For the past two summers, he worked as an engineering aid at Kelsey Hayes Corporation in Brighton, Michigan. Andrew Simms is a senior majoring in Computer Science. His past experience includes three years as a part time employee in the Computing Center Communications group. Bill and Andrew are assisting Rick Schmalgemeier with writing on-line help facilities for NSFNET, working as UMnet postmasters, and creating database applications.

Janise Honeyman recently became a permanent member of the Information Services team as a technical writer. Janise comes to Merit/NSFNET with five years of experience in software documentation. She has worked as writing consultant at University of Michigan Business School, and has taught technical writing, creative writing, and English composition. After completing an M.A. in English, she continued with doctoral studies.

The newly hired coordinator of the Information Services Communications group is Ellen Hoffman. As a staff member of the Merit Technical Support Group, Ellen has been assisting with the NSFNET project since last November but is now permanently assigned to the Merit/NSFNET team. She came to Merit from the University of Michigan Computing Center, where she served as a user consultant and educational coordinator. Ellen has M.A.'s in journalism and anthropology, and has completed coursework towards a doctorate in anthropology with a specialization in archaeology and ethnobotany.

# Link Letters from the backbone

**To: nsfnet-info, nsfnet-info@MERIT.EDU**
**From: steve@note.nsf.gov**
**Subject: I tried to ftp...**

$READ.ME from nis.nsf.net per Link Letter V.1 #5, and got

*$READ.ME: Permission denied*

(Issuing the dir command didn't show anything resembling that filename; is it really there?) Am I missing something? Also, how can I cd back to the top level? On a UNIX system,

*cd ..*

would do the trick - but that doesn't work here.

-s

Steve, thanks for your excellent debugging of our *Link Letter* FTP article. In the sixth paragraph, we say 'you should transfer the $READ.ME file...' but at the time you tried to FTP there were no $READ.ME files on any of the directories. This oversight has been corrected and the $READ.ME file now exists on all of the directories.

The file system under VM is not hierarchical, so the concept of a root or top level directory is not applicable. To get back to the "home" directory for anonymous FTP, a

*cd anonymous*

command will do what you want. Otherwise things worked as advertised in the article.

Since FTP on the IS machine is running under VM on an IBM 4381 mainframe, a couple of other problems have arisen for some users. In the VM/CMS context, filenames starting with a '$' are legal (and force a file to show up at the top of a DIR listing where alpha collation is the default). On some systems, a '$' at the start of a filename may not be legal, or may have a different interpretation. For instance, on a UNIX™ system the leading '$' is interpreted to mean that the following string is the name of a shell variable. Since various systems have different file naming conventions you may need to use the following form of the FTP "get" command to assign a name compatible with your local system when transferring a file:

*get <remote.file> <localfile>*

where "<remote.file>" is the name of the file on the Information Services machine, and "<localfile>" is the name you are transferring the file to on your host. To avoid extra typing, we have put up a READ.ME file that is identical to $READ.ME. We aim to please.

People who FTP to/from IBM mainframe systems need to remember that these machine are EBCDIC-based, and you should <u>not</u> do binary file transfer of a text file. Otherwise, you will end up with an unreadable file on your ASCII system. Do not type <u>BINARY</u> in FTP unless you are transferring executables or other non-text files.

The Link Letter is published by Merit/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of The Link Letter should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions and contributions. IS Staff: Elise Gerich, William Hart, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Andrew Simms, Jim Sweeton.

# Link Letter

The Merit Computer Network

Vol. 1 No. 6 10 June 1988


# GRASPing for information

As many of you know there comes a time during the intense final stages of putting together a large grant proposal when the weary participants move to heights of zaniness otherwise rarely encountered. It was this aura that inspired our name for the vehicle that will ultimately handle all aspects of electronic distribution of information for the Merit/NSFNET team and the NSFNET community — GRASP.

GRASP is a combination of a front-end system (GRANDiose Distributed Application System or GRAND) and a powerful database management system back-end, SPIRES (Stanford Public Information Retrieval System). Together they will provide both standard mechanisms as well as fully-sophisticated functionality to serve our users.

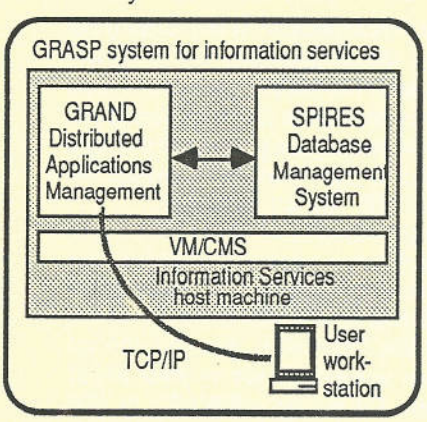


GRAND is a sophisticated server which has been in operation over the vast IBM Corporate Network, VNET (2,000 major computers in some 215 cities throughout the world), for several years. Some of the GRANDiose design considerations exemplify GRAND:

•GRAND supports a variety of communications applications in a "normal electronic environment." Since no single user interface is satisfactory for all users, the individualization of the interface is reasonably easy in the distribution environment

•Authentication of users in a distributed environment is a highly desirable attribute of systems that exchange information. This is particularly critical when some or all of the information is tagged with the poster's name (mail, computer conferencing, for instance). For those applications where such authentication is desirable or mandatory, a GRID (GRAND ID) and associated password are required before the application becomes available to the user.

•Neither central- nor local-server independence is desirable, but complete "cooperation" and concurrence between distributed and central servers is essential. GRAND is quite intelligent in this regard, and there is considerable experience in server-to-server interactions. This distribution of resources among peered servers, with central GRAND management, will be efficient in our backbone environment. We are hoping that some of the backbone sites will be interested in running local GRASP servers.

The GRAND system is based on components called "styles." Styles are both transaction-types as well as the associated software, and GRAND has the ability to define other styles as needed. Fundamental styles include:

•Notification Style (NS) is the transaction that informs the user when a change in an application or entity (file, conference, database) occurs. A common use of NS is for the GRAND computer conferencing application, where the GRAND user (or another GRAND server) wants to know if some other user (perhaps via remote GRAND server) has added comments or a discussion topic.

•Conference Style (CS) is a higher level function than NS, and its job is to manage a particular computer conference model. In the GRAND scheme of things, the concept of a "computer conference" is not a fixed entity, but rather a template for assisting in the creation of an interactive (and perhaps distributed) application that fits the particular needs of the group wishing to communicate.

•User Interface Style (UIS): in a full implementation, UISs are combined to form User Interface Programs. The UISs run in that "user's normal computing environment" that might be anything from VM/CMS to a small microcomputer workstation. This is the front-end of the front-end GRAND system, and coordinates the transactions between the user and the local server, as a two-way stream of styles.

So, for us at Merit/NSFNET Information Systems, GRAND represents a sophisticated set of high- and low-level tools to make our job considerably easier, and to give us a basis for standardizing information-exchange activities in our NSFNET backbone environment. We will be sharing our trials, tribulations, successes, and failures with the GRAND and SPIRES communities, and we look forward to your comments, suggestions, complaints, compliments, and support in our online systems endeavors!

*Note: This is the first of a three-part series on the GRASP System.*

# Link Letter

The Merit/NSFNET Backbone Project

Vol. 1 No. 7 24 June 1988


# A point of view on progress

As many of you know, Saturday evening, June 11, provided an historic moment for the new NSFNET. At that point all thirteen nodes had achieved on-line status and were able to send packets across the backbone. They joined the six test nodes to form, in effect, a sixteen-location, nineteen-node network on schedule. While there remain many tasks to be performed and validated on the remaining schedule of events leading to full production, we should all take a great deal of pride and satisfaction with the accomplishments to date.

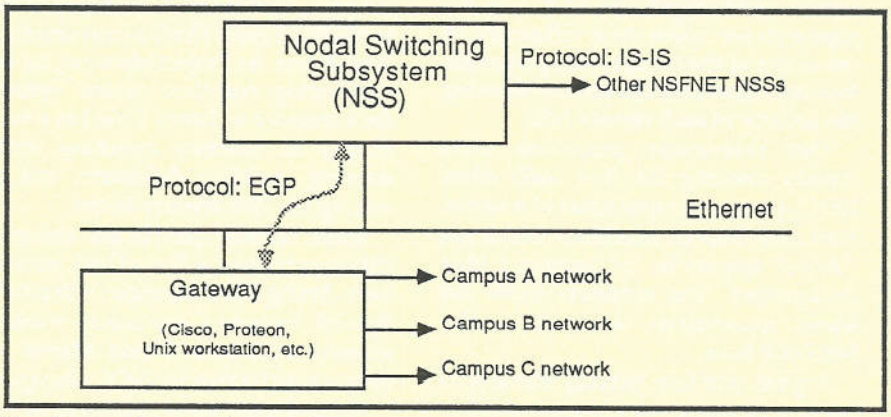


In my view, the single largest ingredient determining the success to date has been TEAMWORK at all levels and across all areas of the NSFNET endeavor. Up until the site folks got their chance, I had been enthused about the teamwork within the formal partnership of MERIT, MCI, and IBM. It's clear now, after observing the extraordinary efforts at the sites—ranging from reconfiguring and swapping equipment to changing cable connectors, to coming to Ann Arbor to input our product plans—the partnership is much bigger than just three corporate names. It's very gratifying. Welcome aboard and thanks!

In the immediate future we'll be applying final software (and some hardware) upgrades, improving the overall network management capability, executing the system test plan (including "break-it"), increasing the amount of regional network traffic, and declaring full production. I believe that if we keep up the outstanding teamwork mentioned above, we'll have a network we can all be proud of within a few weeks.▲

*–Jack Drescher, IBM*

# SPIRES offers sophisticated DBMS and more

In the last issue of the *Link Letter*, we discussed what GRAND is and what GRAND does. This article deals with the joys of SPIRES, and next time we will talk about the combined functionality of GRAND and SPIRES that we call GRASP.

The Stanford Public Information and REtrieval System (SPIRES) is a complete database management system and includes all the tools needed to enter, maintain, search for, and retrieve data. SPIRES finds its roots in the early 70's as a system designed to deal with bibliographic records, and over the years has grown to be a very fine production-level DBMS. The original implementations were developed for the Stanford Timesharing System (STS); by the mid-70's there was a version for the Michigan Terminal System (MTS), and by the early 80's the VM version came into being.

SPIRES now is most popular in the VM environment, and those universities who have SPIRES licenses form the very active SPIRES Consortium. The Consortium sponsors technical workshops in the spring and fall, as well as the week-long Spectrum course on SPIRES programming. This year, the Merit Computer Network and The University of Michigan are pleased to sponsor the Spectrum course in Ann Arbor, October 3–7. Please let us know if you are interested in attending this workshop.

SPIRES implements the hierarchical model for data structuring, though there are very good facilities for doing certain relational operations outside of ad-hoc SPIRES. SPIRES supports variable-length fields ("elements" in SPIRES jargon) and elements may recur as needed. The facilities for element indexing and searching are exceptional—SPIRES databases contain millions of records, yet

*–continued on next page*

# NOC construction complete

Much is happening at the Network Operations Center (NOC): construction is completed, hiring is underway, and the new NOC Manager joined the staff on June 19.

The NOC exists to monitor network traffic and to troubleshoot when problems arise. Traffic for NSFNET and for Merit is monitored here. Located in the Computing Center Building at The University of Michigan, the NOC was designed to use space most efficiently. In the center of the hexagon-shaped room, four operators monitor screens showing the activity of each remote NSS.

Two rear-screen projection units (nearly covering the first wall) allow NOC operators to spot areas of concern. Link problems or machines which are "down" appear in red against a blue background. The monitors show the Merit connections, as well as the NSFNET links.

A glass window looking into a spacious conference room covers another wall. This room will soon have a large television screen to show training and information films. While sitting in the conference room, a visitor (or operator) can view the activity of the network on the projection screens in the central area.

A third window wall looks into the NSFNET machine room where the Ann Arbor NSSs and other equipment reside. The fourth wall looks into the NOC Manager's office and the fifth into the Senior NOC Operator's office. A final solid wall hides a storage room and provides the backdrop for two large whiteboards.

When problems arise on the network, it is crucial that all situations are handled quickly and reliably by the NOC. The first stage of a two-part plan for installing a highly automated information management system at the NOC was completed on June 22. When fully installed, this leading-edge IBM system will record problems as they occur and will then determine if the problem is related to any other known problems, issue an operator request for service when appropriate, and track the progress of the situation until resolved.

A number of staffing changes have already taken place and more will happen within the next few weeks. Dale Johnson, the new NOC Manager, comes to NSFNET from Horiba Instruments in Ann Arbor. As Project Manager, he was developing real-time control systems for automotive clients. Dale has a wide variety of interests including contra dancing, building dulcimers, and restoring his 19th century home.

Some familiar people are assuming different positions. Steve Gold, who has been Acting NOC Manager, will now be Special Projects Coordinator working directly with Dale. David Martin has been named Assistant NOC Manager.

Aaron Landy starts next week as Systems Research Programmer in charge of VM programming.

Currently, the NOC has five operators: Deborah Boyer, Linda Latham-Collins, John Labbe, Kraig Meyer, and Alexis Morgan. Six more operators and two senior operators will soon be appointed.

Activity is brisk here in these last days before the phase-in. As a "hub" for backbone traffic, the NOC is a dynamic place which promises to become a national showcase and state-of-the-art network operations management facility.▲

***SPIRES, cont.***

searching and retrieval of individual records is fast and cheap via the B-tree-structured indexing. In the VM/CMS world, both the full screen-oriented PRISM and line-oriented FOLIO interfaces provide enhanced access to SPIRES databases without the need for any knowledge of SPIRES details.

Our choice of SPIRES, to complement the front-end service of GRAND, was based on our considerable local experience with MTS-SPIRES, and our desire to have a full-function DBMS as a critical part of our network services offerings. In our view, SPIRES is far and away the best software to get the job done!

GRASP, when next we meet.▲

## NSFNET News Briefs . . .

✔On Friday, July 1, the NSFNET Backbone plans to begin the "operational testing" phase of this project. The next *Link Letter* will provide details of this important activity, as operational traffic begins.

✔Beginning July 1, the NOC hours are 8:00 a.m.-12:00 a.m. EDT, Monday-Friday and 8:00 a.m.-6:00 p.m. EDT, Saturday and Sunday. By July 30, the NOC plans 24-hour operations.

✔Doug Elias from Cornell visited Ann Arbor to share information on his experience with gathering statistics for the old backbone. Information Services is working on the structure for supplying these statistics (and others) for NSFNET.

✔An article on monitoring ("Monitoring Data Exchanges between the NSFNET Backbone Network and its Attached Regional Clients") is available in the NSFNET directory on the IS machine by anonymous FTP of the file:

*monitor.mtg-6-88*

The Link Letter is published by the Merit Computer/NSFNET Information Services, 1075 Beal, Ann Arbor, Michigan 48109-2112. Anyone interested in receiving the electronic version of The Link Letter should send a message to: NSFNET-Linkletter-Request@merit.edu. For more information, reach us at NSFNET-info@merit.edu or by calling (313) 764-9430. We welcome suggestions and contributions. IS Staff: Elise Gerich, William Hart, Ellen Hoffman, Janise Honeyman, Laura Kelleher, Rick Schmalgemeier, Andrew Simms, Jim Sweeton.