DOE/ER-0341

U.S. DEPARTMENT OF ENERGY
OFFICE OF ENERGY RESEARCH

ESNET

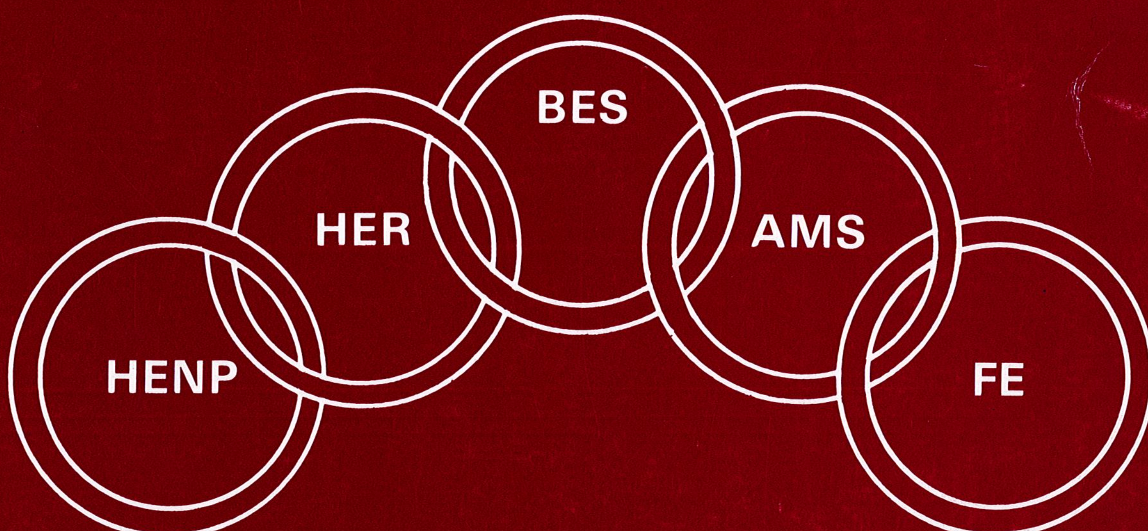

June 1987

Price: Printed Copy A12
Microfiche A01

DOE/ER-0341

U.S. DEPARTMENT OF ENERGY
OFFICE OF ENERGY RESEARCH
SCIENTIFIC COMPUTING STAFF
WASHINGTON, D.C. 20545

ESNET

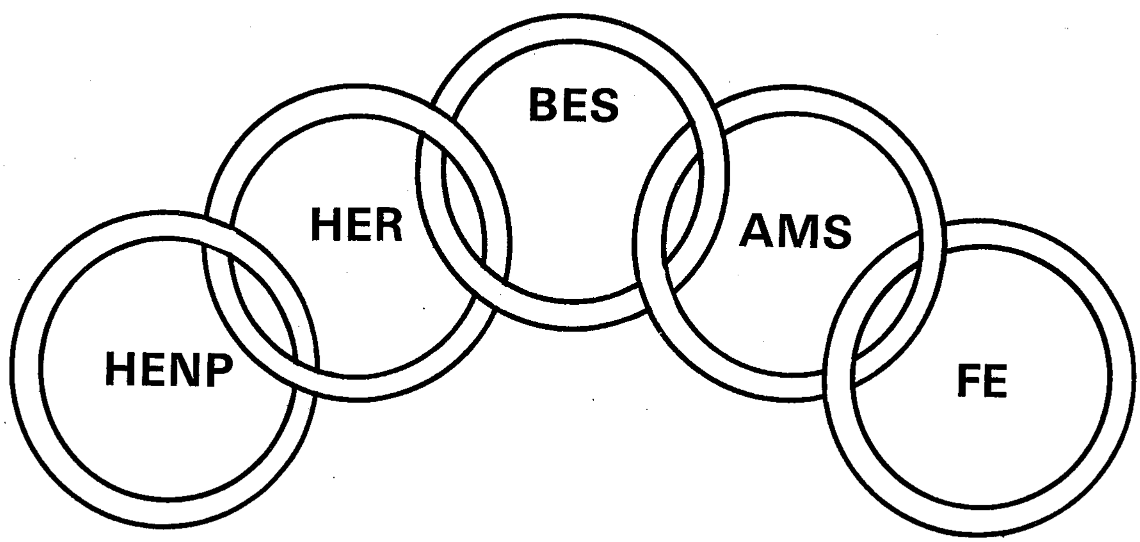

June 1987

DOE F 1325.8
(12-84)

United States Government Department of Energy

memorandum

DATE: July 17, 1987

REPLY TO
ATTN OF: ER-7

SUBJECT: Energy Sciences Network Program Plan

TO: James F. Decker, Acting Director, ER-1
David B. Nelson, Executive Director, ER-2

We are pleased to submit to you this Energy Sciences Network (ESNET) Program Plan in response to your charge to evaluate the computer networking requirement of the Office of Energy Research.

This ESNET program plan codifies and presents the computer networking requirement for all Energy Research programs, analyzes this requirement relative to other Federal research programs, and recommends an implementation plan to address these requirements. This plan is very timely in that there has been a growing awareness of the benefits to be gained in the sciences through the use of computer networks and in that this awareness has even resulted in a Congressional call for a study of computer networking requirements for the U.S. research communities, which has proceeded in parallel to this ESNET study.

We would like to thank the members of the ESNET Steering Committee who have worked diligently to gather programmatic data and requirements and who have assisted in the preparation of this ESNET program plan.

John S. Cavallini

William E. Bostwick
Scientific Computing Staff
Office of Energy Research

INDEX

ESNET PROGRAM PLAN

I. Executive Summary

This document describes the Energy Sciences Network (ESNET) project which was undertaken by the Scientific Computing Staff during fiscal year (FY) 1986 at the direction of the Director, Office of Energy Research (ER). This document serves as the program plan for the ESNET project and is the result of the effort of the cross program Energy Sciences Network Steering Committee. The ESNET Steering Committee has been charged to codify the overall ER computer network requirements, to document and set priorities for computer networking requirements including performance objectives. Further, this committee has been asked to identify future ESNET functional characteristics, to identify research and development needs for the ESNET, to establish ESNET performance objectives and to define the infrastructure necessary to manage and operate the ESNET facilities.

The need for computer networks has been the subject of recent widespread visibility. The success of geographically dispersed scientific collaborations has become increasingly dependent upon the communication facilities which are afforded by computer networking. Networks provide the medium whereby remote scientists can access unique facilities such as accelerators, databases, supercomputers, and even experimental data from spacestations. In addition, computer networks provide the means to integrally connect computer resources at a number of geographically dispersed institutions to support distributed computing capabilities which have shown to be effective in many research applications, especially those supporting large experiments. Computer networks have also been highlighted by several Congressional initiatives, including the National Science Foundation (NSF) supercomputer initiatives, a House of Representatives Science and Technology Committee (HS&T) hearing at the Florida State University (FSU) in June 1985, and the call for a Computer Network Study as part of the FY 1987 NSF authorization bill. The scientific communities within Europe and Japan are currently benefiting from increased visibility and subsequent active Government support.

Computer networks are, in fact, vital to the conduct of modern science, engineering, and technology, and offer great potential to equitably support future collaborations and to provide remote access to scientific facilities from anywhere in the U.S., or in the solar system with the emergence of telescience in experiments on board orbiting spacecraft or spacestations.

To gain these benefits of computer networks, the ESNET concept will include data communications facilities (hardware and software), an internet network architecture plan, and a management infrastructure. The formation of the ESNET Steering Committee was an essential first step in creating the necessary infrastructure.

The initial work by the ESNET Steering Committee resulted in the following recommendations to the Scientific Computing Staff to address the Office of Energy Research computer network needs:

1. A common computer network should be formed in support of all ER programs. This ESNET will be based on the backbone of the existing Magnetic Fusion Energy Network (MFENET) existing Magnetic Fusion Energy Network (MFENET) expanded to include additional ER facilities as shown in Figure VI-2. The fundamental network architecture and direction should follow that proposed in the MFENET Five Year Network Plan (see requirements section) and should include a migration path to support X.25 applications support (see future plans section).

2. ER should provide adequate funding, as a priority, to meet the connectivity and performance requirements for the ESNET as identified in the requirements, foreign access and future plans sections of this plan. In recognition of the open architecture of the ESNET design, ER should encourage and support connectivity to local and campus area networks and to other agency networks, especially the NSFNET for connectivity to university based researchers and facilities.

3. ER should provide adequate support for the ESNET infrastructure. This should include, in addition to facilities and equipment, support for a.) the central computer network research, development, and operation group at the NMFECC, b.) personnel at other major ESNET nodes who would be responsible for applications level network support, local or campus area network connectivity, community of interest network gateway support, and ESNET consulting support, c.) the cross program ESNET Steering Committee to provide direction and to coordinate requirements, d.) support for regular workshops and information exchange in this area, e.) a research and development program to support future applications requirements, and f.) interagency coordination in support of interconnectivity, interoperability, research coordination, and standardization efforts.

4. ER should include detailed information on the ESNET capabilities and facilities to all grantees and contractors as part of the ER contracting process.

II. Introduction

The Scientific Computing Staff manages the supercomputer access program for the Office of Energy Research (ER) to provide high performance computational resources to all research programs that are supported by ER. These programs include High Energy Physics, Nuclear Physics, Materials Sciences, Chemical Sciences, Carbon Dioxide Research, Engineering and Geosciences, Heavy Ion Fusion, Applied Plasma Physics, and Health and Environmental Research. To provide these computational resources, the Scientific Computing Staff manages:

- The National Magnetic Fusion Energy Computer Center (NMFECC) which operates the following supercomputer systems - a Cray 2, a Cray X-MP/22, a Cray 1S and a Cray 1A;

- The Supercomputer Computations Research Institute (SCRI) at the Florida State University (FSU) where research is carried out in the computational sciences related to ER mission areas. The SCRI operates CDC Cyber 205 and ETA-10 supercomputer systems;

- The MFENET which is a nationwide data communications network facility linking over 4,000 ER supercomputer users to these two centers, to locations in Japan and Europe, and to each other from over 100 geographic locations;

- Several advanced computational research centers, such as the Advanced Computational Research Facility (ACRF) at the Argonne National Laboratory (ANL), which operate research/experimental computational resources, e.g., hypercube systems, for use by the research community to support forefront computational research in parallel processing techniques.

During FY 1984, the ER supercomputer access program was significantly expanded in scope. In addition, initiatives to expand supercomputer access were also undertaken by the National Science Foundation, National Aeronautics and space Administration, the National Cancer Institute, the Department of Defense and others. As a result of these initiatives and in recognition of its importance to scientific collaboration support, computer networking emerged as an important issue. [1]

In roughly the same time period, other ER research programs were beginning to join established computer networks or to build networks of their own. For example, both BITNET and ARPANET provided mail and file transfer capabilities for many university research groups and their collaborators at national laboratories. Other university groups found it necessary to lease direct connections to mainframe computers at the remote laboratory where their research activities were concentrated. In the case of High Energy Physics, a private network of leased 9600 baud lines running DECNET grew as university groups required connectivity between their local minicomputers and the facilities at SLAC, Fermilab, Brookhaven, and other HEP labs.

In 1985, a subpanel of the High Energy Physics Advisory Panel (HEPAP) chaired by J. Ballam submitted a report on "Computing for Particle Physics." Chapter 4 of this report, which was entitled "Analysis of Networking Requirements," summarized the situation described in the previous paragraph and recommended the establishment of a HEPNET backbone to provide more effective and efficient networking for the HEP community. This backbone was to consist of high speed (56 Kbps) trunk lines connecting the major HEP laboratories. The report also anticipated that other emerging projects might eventually provide cost effective alternatives to a dedicated HEPNET.

In FY 1985, Dr. Alvin Trivelpiece, Director of ER, charged the Scientific Computing Staff (SCS) to survey the status of and requirements for computer networking throughout all ER programs. This project served as a complementary adjunct to the existing SCS charter for the provision of nationwide access to ER supercomputers. The SCS survey data demonstrated a significant need for improved computer networking to facilitate: 1.) improved access to unique ER scientific facilities, 2.) needed information dissemination among scientific collaborators throughout all ER programs, and 3.) more widespread access to existing supercomputer facilities.

A general purpose computer network architectural concept for ER cross program support emerged for consideration with the redesign of the existing MFENET. [2] This concept was reviewed by an interagency working group in January 1986, and was determined to be a major step toward an interagency internet capability. This internet architectural model (see Figure II-1) offers significant advantages in accessing research facilities and research communities sponsored by other agencies, permits the use of a wide range of vendors' equipment, and interconnects the many university based researchers incorporated into the NSF sponsored networks, such as SURANET.

In response to Dr. Trivelpiece's charge, the Scientific Computing Staff proposed an ESNET concept [3] which recommended the formation of the Energy Sciences Network Steering Committee, offered the new MFENET design and facilities for consideration as the basis for ESNET facilities, proposed an evolutionary operational model, and endorsed a phased approach to addressing long term computer network goals.

During this same period, computer networking has also become a prominent issue in the interagency sphere. The HS&T Committee hearing in Tallahassee, Florida in June 1985, highlighted the importance of computer networks to complement the NSF and other supercomputer access initiatives. Subsequently, the Federal Coordinating Committee for Science, Engineering, and Technology (FCCSET) formed a working group to study this area [4] and Congress mandated a computer network study in the FY 1987 NSF authorization bill.

ER computer network requirements have been developed by using both the Congressional study and the recommendations of the ESNET Steering Committee. These requirements are presented below. Also presented are a summary of the ESNET Steering Committee activities, a statement of the benefits anticipated from the ESNET computer networking project, and a plan of action, including budget requirements for the next 3 years.

ESNET ARCHITECTURAL MODEL

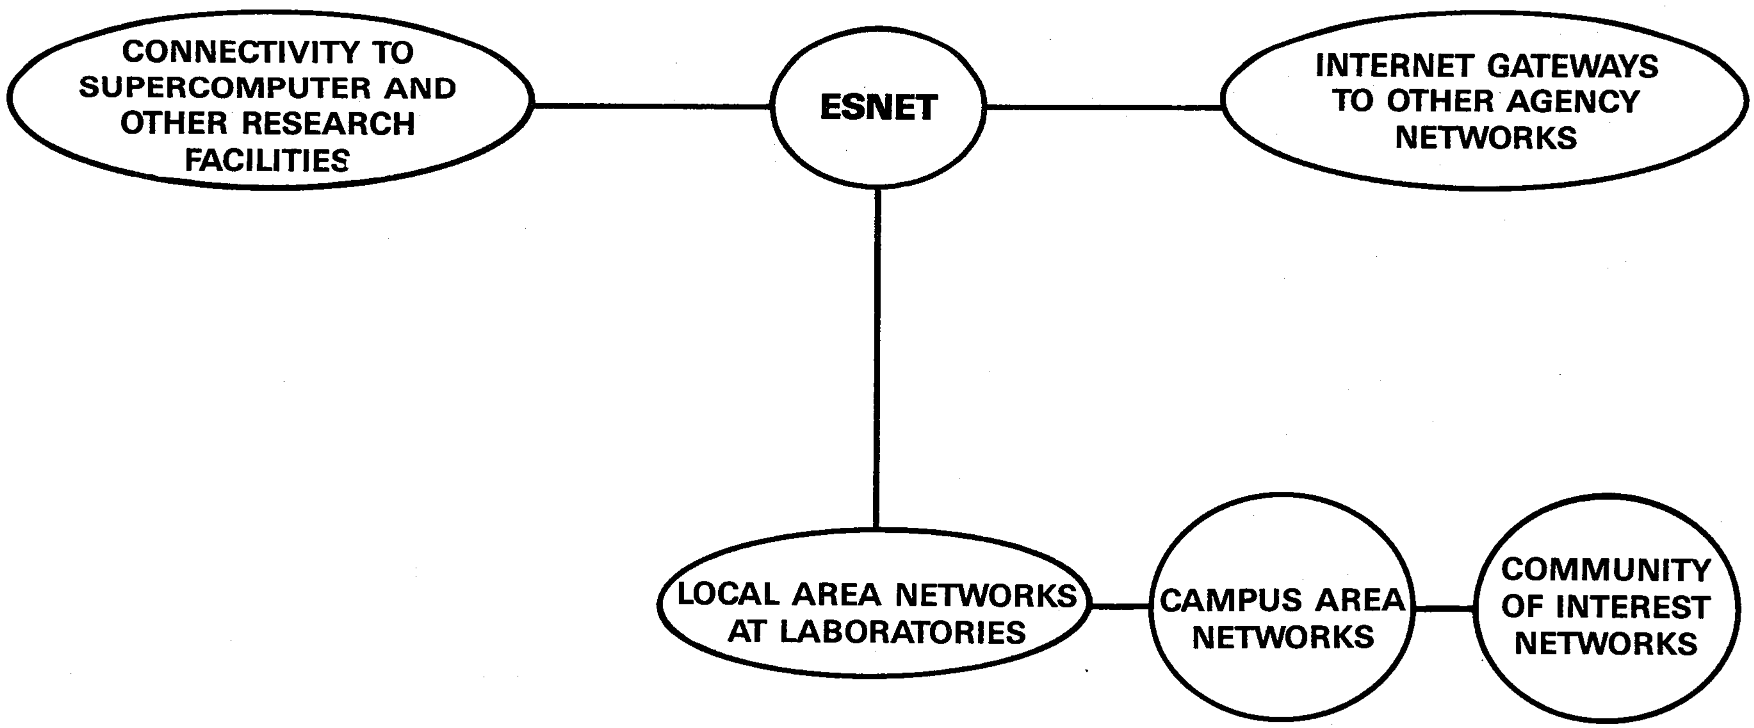

Figure II-1

III. Benefits of Computer Networks to Research

Computer networks offer many benefits to the ER research community. They allow for more cost effective use of scientific research facilities, which are generally large, costly installations such as a synchrotron light source, a particle accelerator, or a supercomputer facility. These installations, because of their high costs, cannot be duplicated at all ER sites. But computer networks, which are properly planned and implemented, have the potential to effectively "bring" these installations to the individual principal investigator. For example, computer networks can support a highly distributed computing environment which is effective in very large experimental collaborations, such as in HEP, where researchers and computing facilities from many sites interact during the staging, execution, and analysis of these experiments. As a result, computer networks have become vital for supporting the research collaborations which utilize these scientific facilities.

Computer networks facilitate the exchange of information upon which research collaborations depend, they enhance the quality and increase the productivity of research projects by providing access to experimental findings/results faster, and they offer more equitable access to these scientific facilities to U.S. researchers without regard to their location or individual funding source. It is also important that computer networks offer these advantages across research program boundaries and therefore, enhance interdisciplinary scientific collaborative efforts as well.

The recent and projected improvements in hardware and software technologies will greatly affect the computing and networking environment available to scientists. Advancements in packet switching, multicasting, fiber optic transmission, and network access protocols will afford more ubiquitous and transparent access to remote research and computational facilities. Developments in distributed operating systems, file systems, and graphics technologies will greatly enhance the functionalities available to the remote researcher. Hence, through proper planning, implementation, and coordination users at out-of-the-way or smaller institutions can reap the same benefits of large scientific facilities as their colleagues at larger institutions or those who are even at the scientific facilities themselves. The lack of such distributed facilities would place ER scientists at a disadvantage compared to their peers worldwide.

For example, a scientist at a small university using a modern, low cost, high performance workstation, if provided with access to a modern, capable computer network, could work in the following fashion. The scientist could have one workstation window opened to a supercomputer application, which is executing over 2,000 miles away, programmed to analyze data for an experiment via another workstation window where this experiment is running at a third installation distant from both the scientist and the supercomputer facility. [5] The scientist's collaborators could also be viewing, by multicast technology, these events at still other geographic locations while electronically communicating with the scientist. Advancements in computer networks are certain to provide qualitative improvements in the ER research programs which they support in much the same way the telephone caused qualitative improvements over the use of letters. In this regard, the most significant benefits of providing qualitative improvements in computer network support for ER research is yet unknown.

IV. ESNET Steering Committee

In October 1986, Dr. Al Trivelpiece, Director of the Office of Energy Research, approved the plan to implement the Energy Sciences Network. [3] The goals for the ESNET concept are to avoid redundant computer network costs via facilities sharing to accommodate the introduction of new technologies in a timely manner, to support multivendor operations, to interface as transparently and seamlessly as possible to other agencies' networks, to provide more equitable access to ER research facilities, to support distributed computing functions and to enhance ER scientific collaborations. To meet these goals, an ESNET steering committee, composed of representatives from all ER programs, was established and was charged to:

- Document, review and prioritize network requirements for all ER programs
- Ensure ESNET goals are achieved without impacting ongoing program requirements
- Establish ESNET computer network performance objectives
- Propose innovative techniques for enhancing ESNET capabilities
- Advise NMFECC network staff
- Identify research needs for addressing network requirements

In mid-October of 1986, William Bostwick of the Los Alamos National Laboratory was detailed to the Scientific Computing Staff on a two year assignment to chair the ESNET Steering Committee. The various ER program areas appointed members from national laboratories, research facilities, and universities to represent their respective research program areas and sites on this steering committee. Those appointed and the programs represented are:

Basic Energy Sciences

Thomas Dunning - Argonne National Laboratory (ANL)
James Davenport - Brookhaven National Laboratory (BNL)

High Energy and Nuclear Physics

Stuart Loken - Lawrence Berkeley Laboratory (LBL)
George Brandenburg - Harvard University
Russell Roberson - Duke University

Fusion Energy

Martin Greenwald - Massachusetts Institute of Technology (MIT)
Richard Hicks - Oak Ridge National Laboratory (ORNL)
Jean-Noel Leboeuf - (ORNL) (effective 6/18/87)

4.8 RECOMMENDATIONS:

A dedicated High Energy Physics Network (HEPNET) should be established to service all Laboratory and University sites.

- This network should be started immediately by funding high speed trunk lines from LBL/SLAC to Fermilab to BNL. A redundant link from BNL to LBL/SLAC may be included in the trunk to improve reliability.
- DOE and NSF funding should permit all university groups to establish a link to one point on this trunk so that those groups can access all other sites.
- A link from BNL to LNS should be established, initially at 9600 Baud.
- A leased line to CERN should be established to provide access to facilities in Europe at the lowest possible cost. The cost of the link to Europe must be independent of the volume of data transmitted on the link.
- A leased line to Japan should be provided when the volume of network traffic is large enough to justify the fixed cost.
- A technical working group should be established to plan and implement HEPNET.
- A permanent staff located at one of the laboratories must be funded to coordinate installation and maintenance of the network. This will require, initally, at least 2 FTE. Beyond the first year, 1 FTE should be sufficient.
- Use must be monitored to assess needs for upgrades.
- Until the new links are working reliably, the existing links should be left intact.

TABLE 4.1

COMPARISON OF NETWORK FEATURES

NETWORK	REMOTE LOG ON	MAIL	FILE TRANSFER	INTERACTIVE MESSAGES	INTERNATIONAL	VENDOR SUPPORT	INTERNETWORK SUPPORT	RELIABILITY
DECNET	YES	YES	TO and FROM	YES	YES X.25 or leased line	YES DEC	BITNET	Requires redundant links for reliable network
BITNET/ EARN	NO	YES	SEND TO only	YES	YES	YES various systems	DECNET	Uses store and forward
X.25/ COLOURED BOOKS	YES	YES	TO and FROM	YES	YES	YES various systems	DECNET one way	Uses retry
DATA SWITCHES	YES	NO	NO	NO	no direct connections	three vendors		Needs continuous links

TABLE 4.2

List of High Energy Physics Computers on Wide-Area Networks

Compiled By Paul Kunz, SLAC

24 December 1984

The following is a survey of computers used by the High Energy Physics community that are attached to wide-area networks. Only computers that support at least mail transfer are considered. Local networks, such as CERNET or DESYNET, are not considered, nor are computers that support only remote logon.

The following networks were considered in the survey:

- BITNET - RSCS network of North American University computer centers
- EARNET - RSCS network of European Research Centers. [BITNET and EARNET are actually one network with different names on each side of Atlantic]
- DECNET - SLAC/LBL DECNET
- JANET - Coloured Books X.25 based network of the U.K.
- INFNET - INFN (Italy) DECNET
- CCNET - Columbia, Carneige-Mellon, Case Western DECNET
- PSSN - Public Packet Switching Networks, e.g. TymNet, TeleNet, only if file/mail transfer is supported.
- ARPANET - DoD network of research centers.
- USENET - UNIX network.

[Network-id in lower case means that connection is not yet made.]

N O R T H   A M E R I C A

Site:/ Dept.	Computer	BITNET/ EARNET	DECNET	OtherNet Contact Person, USERID
Argonne:				
HEP	VAX 11/780	ANLHEP	ANL	
HEP	VAX 11/730		ANLCDF	
Physics	VAX 11/780	ANLPHY		
Center	IBM 3033	ANLOS		
Center	IBM 3033	ANLVM		
LBL & U.C. Berkeley:				
TPC VAX	VAX 11/780		TKYO	
Phys Dept	VAX 11/780	ucbphys?	PHYS	
Phys Dept	VAX 11/ ?		ICS	
Phys Dept	VAX 11/730		TSTBED	
Phys Dept	VAX 11/730		78	
Phys Dept	VAX 11/730		MUTEST	
Brandeis University:				
HEP	VAX 11/780		BRND	
British Columbia:				
Center	Amdahl			MAILNET: UBCG-MTS%MICH-MTS
				B. White
Brookhaven:				
Mailserve	PDP 11	BNL		ARPANET:BNL PEIERLS
Brown University:				
Center	IBM 3081D	BROWNVM		
CalTech:				
Phys VMS	VAX 11/780	SLACVM	CITHEX	PSSN:311021300219 Harvey Newman NEWMAN
Phys UNIX	VAX 11/780			USENET:CITHEP G. Fox
Carneige-Mellon:				
Phys Dept	VAX 11/780			CCNET: CMPHYS
U. Cincinnati:				
Central	Amdahl V8	UCCCMVS		Brian Meadows (PPHYBZM)
Research	IBM 3033N	UCCCVM1		Brian Meadows (PPHYBZM)
U. Colorado:				
Phys Dept	VAX 11/780	bitnet		
Colorado State U.:				
Center	CDC 205	CDC205		(Used by Houston for SSC)
Columbia:				
Nevis	VAX 11/780	cunevis		
Cornell:				
Central	IBM 3081D	CORNELLC		
Central	IBM 4341	CORNELLA		
Wilson lab	VAX 11/780	CRNLNS		Ray Hamlke
Theory	VAX 11/780	CRNLTHRY		
Fermilab:				
Admin.	IBM 4341	FNALVM		Jeff Mack (MAINT)
Front end	VAX 11/780	FNAL	FNAL	Greg Chartrand (GREG)
Front end	VAX 11/780	FNALA	FNALA	"
CD Dev.	VAX 11/780	fnalbsn	BISON	
CDF	VAX 11/780	FNALCDF	CDF	Chris Day, CTDAY
CDF-Soft	VAX 11/780		CDFSFT	"
CDF-Wigwam	VAX 11/730		CDFHRD	"
CDF-Beam	VAX 11/730		CDFNW	"
CDF-Wedge	VAX 11/730		CDFTST	"
ACP	VAX 11/780		BSNDBG	Tom Nash, NASH
Univ of Houston				
Center	AS/9000	UHUPVM1		
Physics	VAX	bitnet		
Harvard University:				
HEP	VAX 11/780	HARVHEP	HUHEPL	
J. Hopkins:				
Phys Dept	VAX 11/780	JHUP	JHU	Jack Serio, SYSTEM

TABLE 4.2

			(dialin)	
U. Ill-Urbana				
Phys Dept	VAX 11/780	UIUCHEPG		Gladdings
Indiana U.:				
	VAX 11/780		IND	
Louisiana State Univ.				
Center	IBM 3081	NSNCC		
U.C. Los Angles:				
Phys Dept	VAX 11/780	uclaphys	UCLA	
U. Maryland:				
Center	Univac 1100	UMD2		
Phys Dept	VAX	bitnet		
MIT				
LNS	VAX 11/780	MITLNS		Robin Verdier, CSC
U. Michigan:				Ian Leedom, LEEDOM
Phys Dept	VAX 11/780		MICH	
	VAX 11/		MICH1	
	VAX 11/		MICH2	
Northeastern Univ.:				
Phys Dept	VAX 11/750	NEUVMS		
NRC - Ottawa:				
Center	IBM 3033	TSSNRCOO		
U. Pennsylvania:				
Physics	IBM 3081D	PENNDRLS		
HEP	VAX	PENNHEP1		
Penn. State				
Central	IBM 3081	PSUVM		
Purdue				
	VAX 11/		PURDUE	
Princeton University:				
Phys Dept	VAX UNIX			usenet ?
U.C. - Riverside:				
	VAX 11/780		UCR	
	VAX		UCR2	
Rockefeller Univ:				
Center	VAX 11/780		ROCKVAX	
U.C. San Diego				
Phys Dept	VAX		SDPH1	
U.C.- Santa Barbara:				
Phys Dept	VAX 11/780	bitnet	UCSB	Steve Yellin, SPYPO9

Stanford University:				
Central	IBM 3081K	STANFORD		
SLAC:				
Central	IBM 3081K	SLACVM		Les Cottrell, COTTRELL
MAC	VAX 11/780	SLACMAC	MAC	Dave Wiser, DEWPO6
CB	VAX 11/780	SLACCB	CB	Didier Bisset, DIDIER
MarkIII	VAX 11/780	SLACMK3	MK3	Dennis Wycinski
MarkII	VAX 11/780	SLACMKII	MKII	Andy Lankford
ASP	VAX 11/750	SLACASP	ASP	Gabor Bartha
TBF	VAX 11/750	SLACTBF	TBF	Dave Sherden (DJS)
SLD	VAX 11/780	SLACSLD	SLD	Dave Sherden (DJS)
End Stn A	VAX 11/780		ESA	Zen Salata
HRS	VAX 11/780	slachrs	HRS	Jim Schareth
TPC	VAX 11/780	slactpcs	TPCS	Ed Whipple
TPC	VAX 11/780		TPCT	Ed Whipple
Two Gamma	VAX 11/780	SLACTWGM	TWGM	Ed Miller
Two Gamma	VAX 11/750	SLACUCSD	UCSD	Ed Miller
Two Gamma	VAX 11/750	SLACNIKH	NIKHEF	Ed Miller
Two Gamma	VAX 11/780	SLACUCD	UCD	Dave Pellett
PEP C.R.	VAX 11/780		PCR	Tony Gromme
SLC	VAX 11/780	SLACSLC	SLC	Nan Phiney
MCC	VAX 11/780		MCC	Nan Phiney
SPEAR	VAX 11/780		SPEAR	Mike Sullenberger
StonyBrook:				
HEP	VAX			DECNET: to where?
Syracuse University:				
HEP	VAX 11/782	SUHEP		
U. Tennessee:				
Central	IBM 4341	UTKVM1		Jim Brau, PA87178
Texas A&M:				
Phys Dept	VAX	TAMVXPHY		
Vanderbilt U.:				
Phys Dept	VAX 11/780	VANDVMS1		Frank Kyle, KYLE
VPI:				
Center		VPIVAX3		Luke Mo
U. Wisconsin:				
PSL-A	VAX 11/780	WISCPSLA	FERMI8	
PSL-B	VAX 11/780	WISCPSLB	WISC	
Yale University:				
Center	IBM 4341	YALEVM		
Center	VAX	YALEVAX5		

E U R O P E

[The JANET nodes are based on note written by James Hutton, PAPO5/54-1, Rev 5.2, DCC 84-5, 16 November 1984.]

Site: Dept	Computer	BITNET/ EARNET	JANET	OtherNet Contact
Germany				
Univ. Bonn				
Central	IBM 3081	DBNRHRZ2		
DESY:				
Central	IBM 3081D	DHHDESY3		
Mark-J	VAX 11/780			PSSN:2624540009306
				Coloured Books
				Richard Mount, MOUNT
Tasso	VAX 11/750		DYVB[*]	K. Gather
[* Restricted to U.K. users only, due to be moved to PPSN]				
Heidelberg:				
HEP	IBM 4341	DHDIHEP1		
MPI-Munich:				
Central	IBM 4341	DMOMPI11		
Netherlands				
Kath. Uni Nijmegon:				
Central	NAS9040	HNYURC11		
Israel				
Weizmann Institute:				
Central	IBM 3081D	WEIZMANN		
Tel Aviv:				
Central	IBM 4341	TAUNIVM		
Central	CDC 170	TAUNNOS		
Technion University:				
Central	IBM 3081D	TECHNION		
Italy				
Bari:				
CSATA	IBM 4341	IBACSATA		
INFN	VAX			INFNET:VAXBA
Bologna:				
Center	IBM 4341	ICINECA		
INFN	VAX	earnet?		INFNET:VXBO(12), Sarterelli
gateway	PDP11/70	IBOINFN		INFNET:TOCNAF(11),Gateway
Frascati:				
	VAX			INFNET:VAXLNF(5)
Genova				
	VAX			INFNET:VAXGE(19)
Milan:				
	VAX			INFNET:VAXMI(22)
Padova:				
	VAX			INFNET:VXFPD
Pisa:				
CNUCE	IBM 3033	ICNUCEVM		
CNUCE	IBM 370/168	ICNUCEVS		
INFN	IBM 4341	IPIINFN		
INFN	VAX 11/750	IPIVAXIN		INFNET:VAXPI(20)
Rome:				
	VAX 11/780			INFNET:VXROM(6) Valente
	?			INFNET:CBROM(7)
	?			INFNET:OGROM(8)
Trieste:				
	VAX			INFNET:VAXTS(21), Liello
Switzerland				
CERN:				
Central	IBM WYLBUR	CEN		USENET:CERNIBM
Central	IBM VM	CERNVM		
3081/E	IBM 4361	cernvme		
DD Dev	VAX 11/780	CRVXDEV	CPVA[*]	INFNET:VXDEV(60)
OPAL	VAX 11/780		CPVB[*]	N. Gee
Omega	VAX 11/780		CPVC[*]	W. Carena-Bozzoli
Merlin	VAX 11/780		CPVD[*]	D. Myers
Delphi	VAX		CPVE[*]	G. Mornacchi
Gateway	?			INFNET:CERNGW (10)
OPAL	VAX			INFNET:VXHYP(48)
LEP DB2	VAX			INFNET:VXLDB2(37)
Gift Proj	VAX			INFNET:VXGIFT(40)
L3	VAX 11/750			INFNET:VAXL3(41)
Aleph Dev	VAX	crvxalfb		INFNET:VAXALFB(54)
Aleph beam	VAX		CPVF[*]	INFNET:VAXALBM(55)
Aleph TPC	VAX	CRVXALTP		INFNET:VAXALTP(56)
Unix	VAX	earnet		USENET:CERNVAX
[* restricted to U.K. users only, due to move CERN X.25 network in January 1985]				
Universite de Geneve				
	VAX	CGEUGE51		
United Kingdom				
Birmingham U:				
	IBM 4341		BHIA	L. Lowe
Bristol Univ:				
	VAX 11/750		BRVA	J. Alcock
Cambridge Univ:				

Site	Machine	Node		Contact
	VAX 11/780	CAVA		R Ansorge
Daresbury Laboratory:				
Mail Server	GEC	DLGM		
Edinburgh Univ:				
	VAX 11/750	EDVA		D. Candlin
Glasgow Univ:				
Nat Phil.	IBM 4361	GWIA		A. Conway
Imperial College:				
	VAX 11/780	ZIVA		R. Beuselinck
	IBM 4341	ZIIA		R. Campbell
Lancaster Univ:				
	VAX 11/750	LAVA		R. Henderson
Liverpool Univ:				
	GEC 4085	LLGA		M. Houlden
	IBM 4331	LLIA		
Manchester Univ:				
	GEC 4090	MAGA		R. Hughes-Jones
Oxford University:				
Nucl Phys	VAX 11/780	XXVA		J. Macallister
NP	DEC 10	XXDA		W. Black (BLACK)
Queen Mary College:				
	VAX 11/750	ZMVA		P. Kyberd
Rutherford laboratory:				
Central	IBM 3081(VM)	RLIB	rlvm370	R. Maybury
Central	GEC 4090	RLGB		K. Duffey
HEP	VAX 11/780	RLVB		M. Waters
Library	Prime	RLPE		
Sheffield Univ:				
	GEC 4085	SHGA		C. Wells.
Southampton:				
	GEC 4070	SNGA		M. Counihan
Surrey Univ:				
	Prime 550	SYPE		
Sussex Univ:				
	Prime 550	SVPA		
University College, London:				
	GEC 4085	ZUGA		J. Conboy

Gateways between Networks:

| From: \ To: BITNET DECNET JANET INFNET CCNET ARPANET |

Gateways between Networks:

From: \ To:	BITNET	DECNET	JANET	INFNET	CCNET	ARPANET
BITNET	-	SLACCB	ukacrl	iboinfn	CUVMB	UCBJADE
DECNET	CB	-				
JANET	????		-	gift		ZUXA[*]
INFNET	????		vxgift	-		
CCNET	CUVMB				-	
ARPANET	BERKELEY		UCL-CS			-

[Note: Lower case indicates not yet operational or planned]
[BITNET and EARNET are considered the same network]
[* restricted to registered ARPANET users]

Table 4.3
Typical Transfer Rates

Desired	Mid 80's to 1990's	Early 1990's
File transfer		
Switched	9.6	56 kb
pt. to pt.	56 - 224	>1 mb
Broadcast	9.6 → 3 mbytes/hr.	

Example	Size	Goal	Time
LEP3 Small MC File 40 Mbytes	5×10^8 bits	56 kb	2 hrs
any 100k Fortran lines executable images	6 Mbytes	9.6 kb	2 hrs
TPC calibration file	1 - 3 Mbytes	9.6 kb	20 - 80 min
any graphics 3D image	10^6 bytes	56 kb	2 min
[LEP ∿ 10 tapes	1 Gbytes	1.544 Mbits	2 hrs]

BITNET

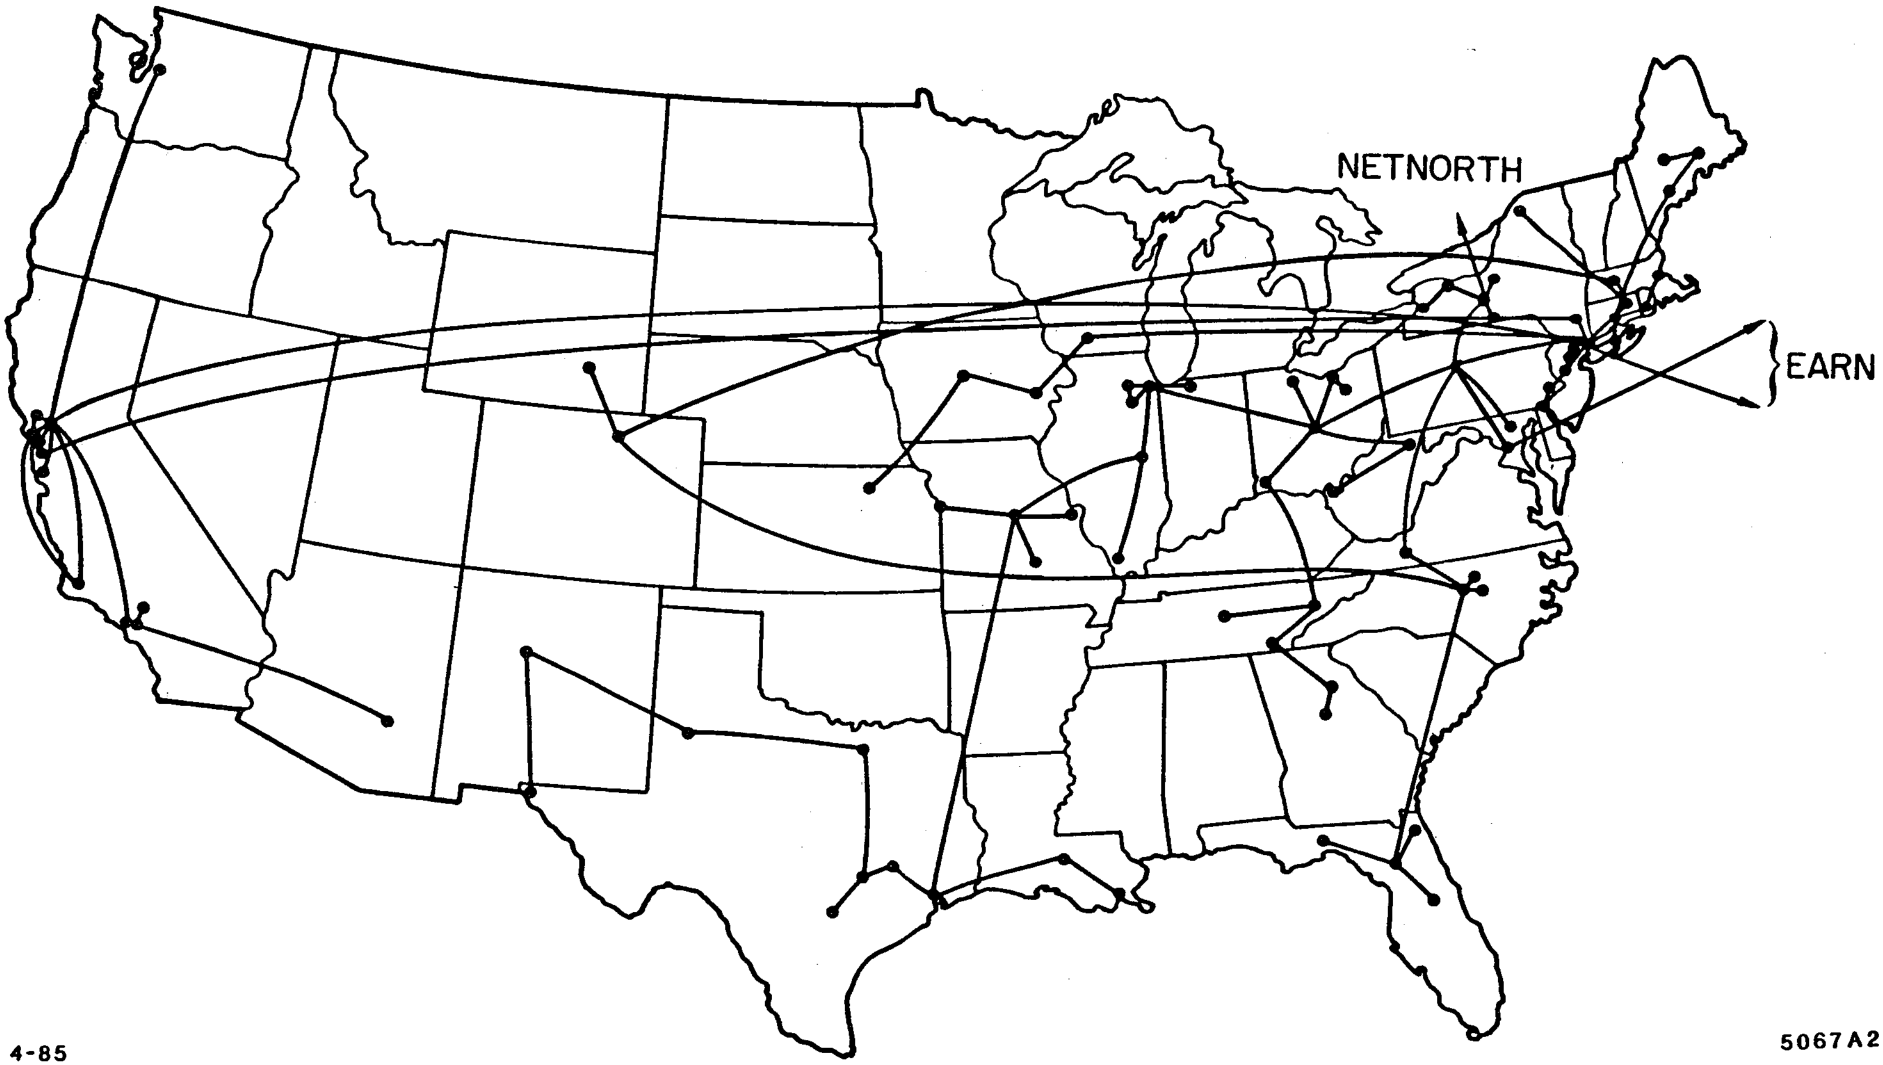

4-85

5067A2

FIGURE 4.1

HEP LEASED LINES FOR TERMINALS

4-85

5067A3

FIGURE 4.2

SLAC/LBL DECNET

4-85

5067A4

FIGURE 4.3

HEPNET PHASE I

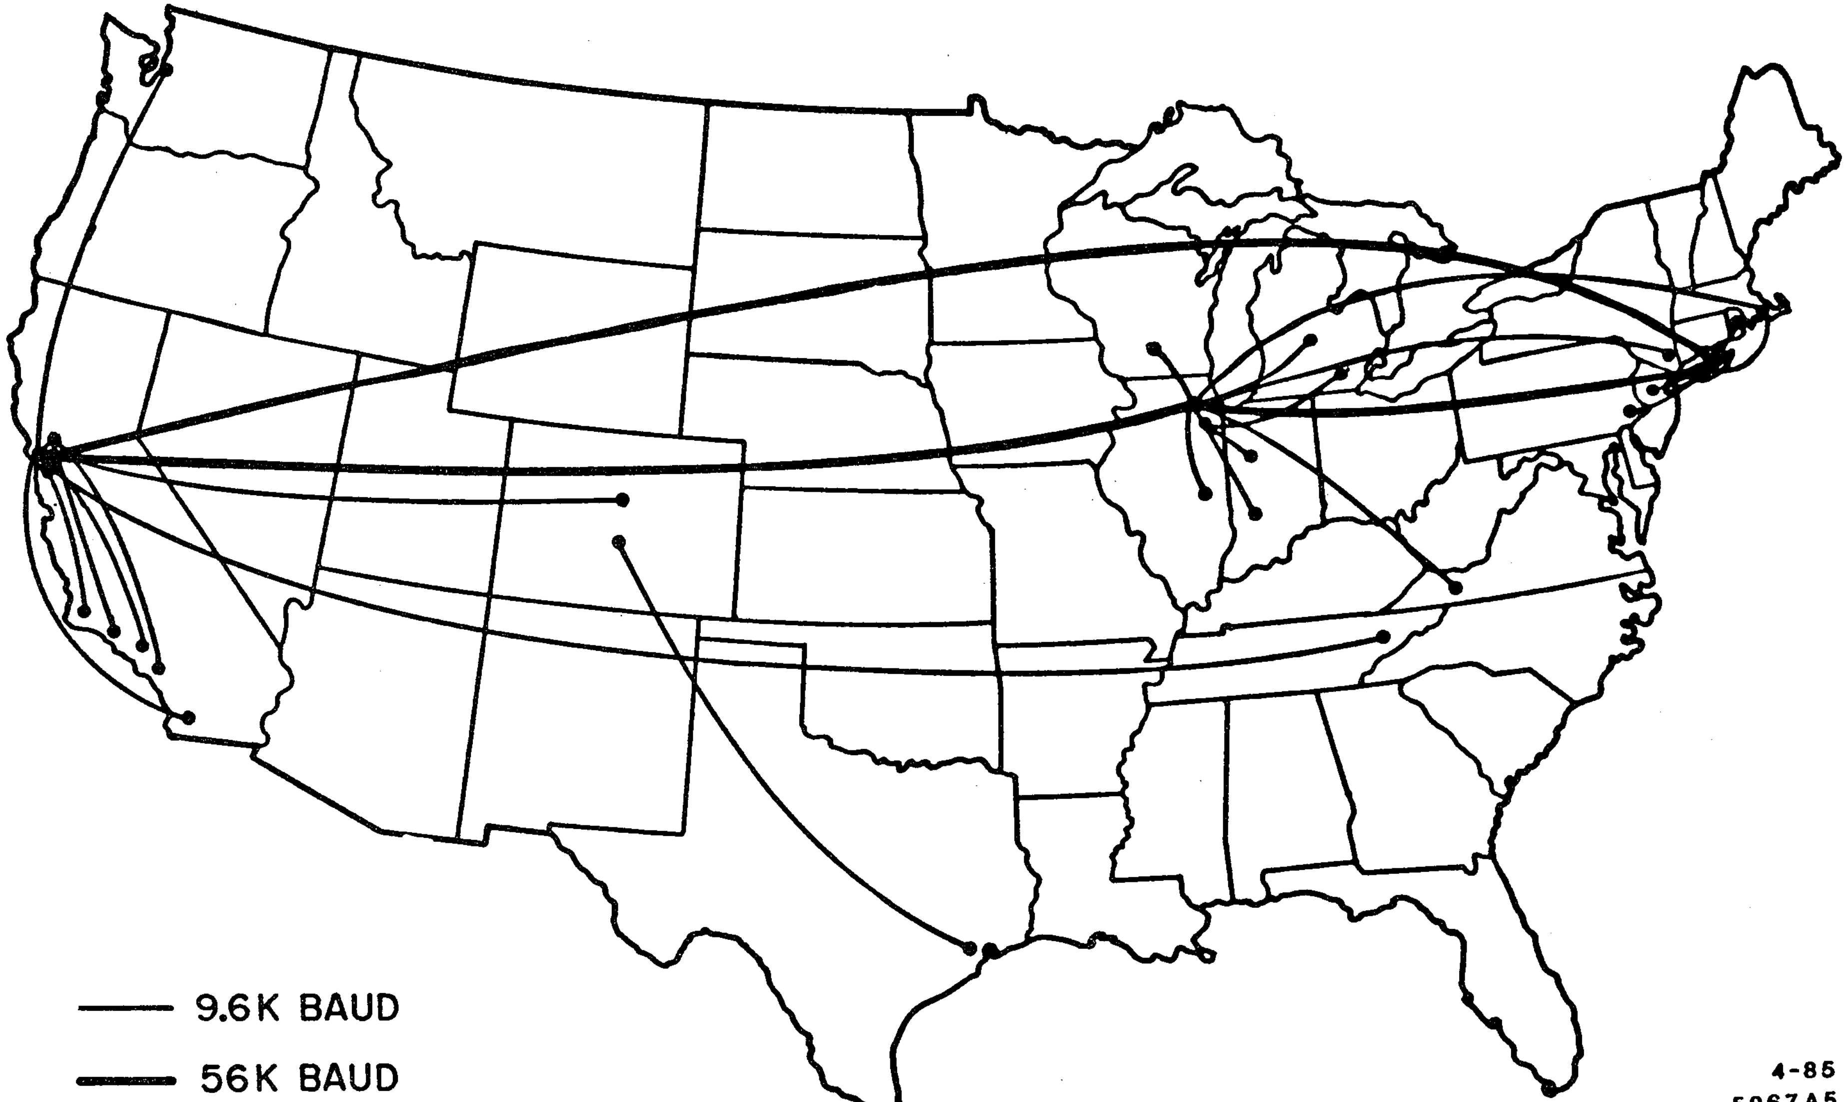

FIGURE 4.4

SLAC MEMORANDUM

19 September 1986

TO: HEPNET Technical Committee

FROM: Howard Davies

SUBJECT: Equipment Requirements at ESNET HEP Sites

---

1. Introduction

At the last meeting of the HTCC on 21 May 1986, an outline proposal of the facilities to be offered by ESNET at the principal node sites was put forward. The proposal included, at least as an interim measure, the provision of an X25 carrier service between ESNET nodes which would allow HEP sites to mount services which have a very high priority in the physics community, namely high speed DECnet file transfers and terminal access to large mainframes installed at the major laboratories

Many details of the X25 service over ESNET remain to be defined. On the basis of the outline proposal made in May, this note describes the kind of equipment which will be required at the major HEP nodes, lists a number of points which must be confirmed with the ESNET development group before detailed planning can proceed, and makes a rough estimate of equipment costs.

2. ESNET Services: a Functional Description

The services to be offered on the ESNET equipment installed initially at the principal HEP laboratories and the ways in which this ESNET equipment will be linked to local computing and communications equipment are shown symbolically in Figure 1.

Briefly, ESNET will offer a number of services based on the use of the TCP/IP protocols via equipment (labelled ESNET Controller in Figure 1) which will be connected to the node site's Ethernet, thus making ESNET services available, at least in principle, to most other computing equipment on the site. Both the satellite channel and the landline (for terminal traffic) connecting one ESNET node to others will have bandwidths of at least 56 kbit/sec. As a separate service, but one with a higher priority for the HEP community in the short term, ESNET will also carry X25 traffic between pairs of node sites. The nominal throughput available for X25 traffic will also be 56 kbit/sec and some form of multiplexing arrangement will be provided to handle both the X25-only and the "normal" ESNET traffic on the same long distance link to the next ESNET node. Because more than one device will need to make use of the X25 link, an X25 switch will be required at each node site. The switch will be connected by a 56kbit/sec line to at least one local host computer, possibly a DECnet machine. Further connections may be made at slower speeds, say 9.6 kbit/sec, to other local host computers capable of handling file transfers over X25. At least one 9.6 kbit/sec link will be required between the switch and an X25 PAD which will handle terminal connections both from local terminals to hosts at remote sites and (in reverse PAD mode) from terminals at remote sites to local hosts. If the sum of the number of local terminal connections and the number of reverse PAD connections offered exceeds 16, more than one PAD (and line to the X25 switch) is likely to be required.

In practice, two or more of the functions described above may be combined in a single piece of equipment. For example, one device might serve as ESNET controller and as the multiplexor for TCP/IP and X25 traffic. It is assumed that ESNET will support simultaneous X25 calls from one node to several remote nodes and will therefore have X25 switching capability. If the ESNET node equipment is capable of X25 switching, a separate switch would not be necessary. Similarly, both X25 switching and PAD functions might be combined in a single piece of equipment. More information on the plans of the ESNET development group is required before the functional requirements listed above can be translated into a detailed equipment specification. (A list of questions for the ESNET development group is attached as Annex A.) In the meantime, an equipment layout in which each function described in Figure 1 is supplied by a separate piece of equipment will serve as the basis for the next stage of planning.

3. Performance Requirements

The most important performance parameter for any kind of X25 switching equipment is its overall throughput measured in packets per second. As a first approximation, the central processor time required to deal with an incoming packet (read the data from the line, interpret the header, decide where the packet has to go, output the data on another line) is independent of packet size. For HDLC transfers, line handling is looked after by special purpose chips and contributes little to the central processor (CP) load. The CP overhead per packet therefore determines throughput capacity. In the case of a switch configured as in Figure 1, traffic on the two 56 kbit/sec lines will form a large proportion of the load and the required switch capacity is approximately equal to that required to handle this traffic.

Equations describing packet rates and data throughput in terms of line speed, packet size, numbers of calls etc are given in Annex B. The basic approach is to ask what level of service can be provided if equipment similar to that described in Figure 1 is available (and the implied transmission bandwidth constraints are accepted). Because of the large number of variables involved, it is necessary to make some fairly arbitrary assumptions about the values of some of them in order to calculate the capacity required for ESNET node equipment. For the purpose of this note, the following assumptions are made:

(a) Effective line utilization—ie the number of useful bits, including those making up HDLC and X25 Level 3 packet headers as well as user data, divided by the nominal capacity of the line—is 100%. (In practice, bit stuffing and the transmission of extra HDLC flag frames will reduce the line utilization below 100%, but only by a small amount. Although such high line utilizations are routinely achieved at speeds up to 9600 bit/sec, it is not certain that current HDLC interfacing equipment can support lines which are an order of magnitude faster equally well. This assumption is justified as representing the "worst case" load on the switch.)

(b) The total load is symmetrical and evenly balanced over the two directions of transfer. Since channels used for X25 traffic are full duplex, the actual amount of data transferred—summed over both directions—is then equal to twice the nominal line capacity, ie 112 kbit/sec for the high speed ESNET lines.

(c) All terminal activity is in screen handling mode. Each character entered at a keyboard is sent in its own packet to the host computer which replies with an echo character, also transmitted one per packet. (Some terminals can generate two or three characters as a result of a single keystroke, eg to indicate cursor movement, and receive more than one character in reply. If these groups of characters are transmitted within a single packet, the effect on the results of the calculations is small.)

(d) The average rate at which each user inputs a character at a terminal keyboard is 1.5 per sec and the average number of output characters generated (in addition to echo characters) per terminal is 170 per sec. (These values are based on measurements of terminal activity of users accessing SLACVM. Figures for the busiest second of the day have been used. Rates averaged over the busiest minute and busiest hour of the day are about 60% and 25% respectively of these values).

With these assumptions, data rates, packet rates and the number of terminal calls supported are shown graphically in Figure 2 as a function of the ratio of the load from file transfers to the total load (from file transfers and from terminal traffic). Figure 2 shows two sets of plots corresponding to maximum packet sizes of 128 and 256 bytes, the most likely values. Note that the File Transfer/Total load ratio is defined in terms of the number of packets generated by each type of traffic; the load ratio in terms of bytes transmitted is different except at the end points of the graphs.

The load ratio arising in practice is likely to depend on whether the network will be used regularly for bulk data transfers, eg to move large amounts of experimental data from the laboratories to physicists' home sites. If so, the load ratio will approach the high end of the range. If, on the other hand, file transfers are limited to the exchange of mail, programs, documentation and occasional samples of experimental data, the requirement to handle terminal traffic at peak times will predominate and the configuration should be planned to handle a load ratio in the lower half of the range.

		Packet Size (bytes)	
		128	256
File Transfer Traffic	byte/sec	7200	9100
	packet/sec	56	36
Terminal Traffic	Number of Terminals	30	23
	byte/sec	5250	3900
	packet/sec	131	83
Total User Traffic	byte/sec	12450	13000
	packet/sec	187	119

Table 1: Switch Traffic with a Load Ratio of 0.3

The assumptions and the corresponding Figure 2 values represent a set of scenarios which can be used to derive a switch specification. As an example, if the assumptions in (a) to (d) above are considered reasonable and the further assumption is made that, at peak load, 70% of the traffic results from terminal activity and 30% from file transfers (ie the load ratio is 0.3), the parameter values are as shown in Table 2 and the specification for the switch might be as follows:

— Support for at least two 56 kbit/sec lines.

— Support for at least four 9600 bit/sec lines in addition.

— Overall throughput of at least 200 packet/sec (based on a packet size of 128 byte).

Similar calculations can be carried out for an X25 PAD and lead to a possible specification of PAD capacity as follows:

— Support for at least one 9600 bit/sec line.

— Support for at least 16 terminals in any combination of PAD and reverse PAD mode.

— Overall throughput of at least 50 packet/sec.

4. Equipment Costs

It is assumed that the ESNET Controller and Multiplexor will be specified, budgeted for and supplied by the ESNET Development group and that its cost need not be discussed by the HTCC. However, the proposed ESNET services include only a carrier service for X25 traffic and the HEP community will have to seek additional funds for the purchase of the new equipment required to make use of the X25 carrier facility. The equipment to be funded consists of the X25 switch, PAD(s), host interface(s) and cables at each ESNET physics node. Unless it is provided as part of the standard ESNET service, some diagnostic and monitoring equipment may also be required.

A budgetary estimate for the cost of an X25 switch as specified in section 3 above is $8000 (based on the price at current exchange rates of a Camtec *SwitchPAD*) A typical 16-port PAD, the Micom *Box Type 2* with *X25 PAD Featurepak*, costs $4300 but a single device would not meet the specification given above—at least two such PADs would be required to provide a throughput capacity of 50 packet/sec and five or six might be needed in the scenario defined by Table 1. *

* Micom quote a capacity of 15 packet/sec for their *Box Type 2* but rates three times higher than this have

Annex A

Questions for the ESNET Development Group

1. Is the Functional Description of the ESNET X25 services given in Section 2 of the main text accurate?
2. Will the X25 (1980) or X25 (1984) protocol, or both, be supported?
3. Will a maximum window size of 128 be supported for X25 (Level 3) packets?
4. What is the maximum User Data Field length that will be supported?
5. Will Flow Control Parameter negotiation be supported (it may be useful to use different values of these parameters for calls of different types, eg file transfer or terminal handling, calls passing via a satellite link or via landline)?
6. How will routing of file transfer calls (via satellite) and terminal handling calls (via landline) between a single pair of source and destination machines be managed? Will machines at user sites be required to have two (or more) X25 addresses in order to support traffic of different types?

$$U_t = \frac{p(1-R)(2c_i+c_o)}{k_2} L_{tot}$$

$$P_f = \frac{Rk_1}{k_2} L_{tot}$$

$$P_t = \frac{k_1(1-R)}{k_2} L_{tot}$$

$$P_{tot} = \frac{k_1}{k_2} L_{tot}$$

where

$$k_1 = (2c_i p + c_o)$$

and

$$k_2 = \left(2c_i R p^2 + (18c_i R + 2c_i + 20 - 20R + c_o)p + 10c_o\right)$$

4. Substituting the parameter values given in the main text, namely $c_i = 1.5$, and $c_o = 170$,

$$k_1 = 3p + 170$$

$$k_2 = 3Rp^2 + (7R + 193)p + 1700$$

5. For a full duplex 56 kbit/sec line with the full bi-directional capacity being used to transmit useful data bytes (packet headers and user data),

$$L_{tot} = (2 \times 56000)/8 = 14000$$

TENTATIVE - STU LOKEN

FOREIGN REQUIREMENTS

Program High Energy Physics

FROM	TO	APPROXIMATE DATE	TYPE	TECHNICAL CONTACT
MIT	CERN (Geneva, Switzerland)	Installed	land (19.6K)	H. Newman (Cal Tech.)
Fermilab	CERN	Summer 1987	satellite (56K)	B. Carpenter (CERN)
LBL	Japan (KEK)	Immediate (short term)	satellite (9.6K)	Karita (KEK)
LBL	KEK	Summer 1988	satellite (64K)	Karita (KEK)
Fermilab	DESY (Hamburg, FRG)	1988-1989	satellite (9.6 - 56)	H. Hofmann (DESY)
Fermilab	Italy	1987	satellite ? (9.6) paid by Italy	E. Valente (INFN Rome)

COMPUTER NETWORK NEEDS - AFRD

Summary

The Accelerator and Fusion Research Division at LBL contains a variety of projects with diverse computing needs. All groups use computer networks to send electronic mail to collaborators at other labs in the U.S. and Europe. Mail messages include both information and source code. Almost all groups require access to supercomputers off-site at the NMFECC at Lawrence Livermore National Laboratory. This entails both access for terminals, and file transfer capabilities. Source code and data are sent to the supercomputer, and text, data, and graphics files are transferred back to LBL to be printed or, occasionally in the case of text, edited. At present one group, the Center for X-ray Optics, receives data from experiments at other laboratories via Decnet, and uses the same network to both use Vax computers at other sites and allow access to the LBL Vaxes from those sites. Tables 1 and 2 below give a more detailed and quantitative accounting of present network use. Table 1 shows the number of terminal connect hours, file transfers, and electronic mail messages used by the division per month. Table 2 lists the sites to which these connections occur, and the networks presently used.

Over the next 3-5 years we estimate that the present patterns and amount of usage will be maintained, with a slow increase in the amount of electronic mail.

Table I

Group	Messages/mo.	Terminal Hrs./mo.	Files/mo.
Bevalac	6/year	0	0
Center for X-Ray Optics	20	2	2
Exploratory Projects	100	.22	29
Heavy Ion Fusion	28	85	130
Heavy Ion Linear Accelerator	16	85	25
Magnetic Fusion Energy	3	265	1000
Superconducting Magnet Design	50	160	140
Two-Beam Accelerator	4	225	60
Totals	221.5	822.22	1386

Table 2

Group	Connection to:	Network
Bevalac	other labs in U.S. & Europe	Bitnet
Center for X-Ray Optics	Wisconsin, BNL Germany, Israel	Decnet, Bitnet
Exploratory Projects	ANL, BNL, FNAL, CERN, Frascati, Bessy (Berlin), Daresbury (UK), LANL	Decnet, Bitnet, Arpanet, MFENet
Heavy Ion Fusion	CERN, Cornell, LASL, DOE, NRL, LLNL, MFECC	Bitnet, Arpanet, MFENet
Heavy Ion Linear Accelerator	MFECC, GSI	MFEnet, Bitnet
Magnetic Fusion Energy	MFECC, GSI	MFEnet, Bitnet
Superconducting Magnet Design	BNL, FNAL, misc. other labs, MFECC	Bitnet, Arpanet, MFEnet
Two-Beam Accelerator	MFECC	MFEnet

Functional Characteristics of ESNet: Nuclear Physics

The ESNet functionality needed by the nuclear physics community is not basically different from that required by the other ER-subfields. The programmatic requirements for nuclear physics will be met if ESNet has the characteristics spelled out in the Dunning, Greenwald and Loken report (in its final form) and if *every* user has the opportunity to connect to the network.

The hardware connectivity requirements appear, at first sight, to be different for the two branches of nuclear physics (low- and intermediate-energy nuclear physics) for the following reasons:

1) Most low-energy nuclear physics (experimental and theory) is carried out at university sites by university personnel.

2) Intermediate-energy nuclear physics is most often carried out at national laboratory facilties, such as LAMPF at LANL, by both laboratory personnel and by university user groups.

The main needs of low-energy (LE) nuclear physicists will be a) interactive and batch service on remote supercomputers, b) E-mail service, and c) file transfer of documents. The university user groups of intermediate energy (IE) facilities will have the same needs plus d) interactive and batch service on the computers located at the IE facilities and e) file transfer and manipulation of large data bases. So, in fact, the connectivity requirements are really not very different for IE and LE nuclear physics.

Each group, whether it is at a national lab or at a university, will need a link to ESNet with adequate speeds (see Dunning, Greenwald and Loken report). Therefore, the central question is whether (1) ESNet is to include the backbone plus the links to the various universities and sites at national labs, or (2) is ESNet only the backbone. If model (1) is used then the question is how are links going to be funded. Who is going to fund the connection of a university group to one of the ESNet nodes? In the case of HEP, the university groups have used their own funds to pay of lines to Fermi Lab, etc, but many of the smaller IE and LE groups do not have the necessary funds to do this. If model (2) is assumed, then the only hope for ER funded university groups is to be able to reach ESNet via the NSF regional networks. The interconnections must be transparent, from a user standpoint must be "flat", and must provide the functionality discussed in the Dunning, Greenwald and Loken report. For this model, the question of chargeback will have to be settled by DOE and NSF at the headquarters level.

FOREIGN REQUIREMENTS
Nuclear Physics

From	To	Date	Type	Contact
Carnegie Mellon	SPS	?	S,T	G.Franklin
	LEAR	?	S,T	(P. Barnes)
BNL	CERN	>12/88	S	M. Levine
	Heidelberg	>12/88	S	(Ole Hansen)
	KEK	12/89	S,T	
U. of Houston	CERN	?	S	Larry Pinsky
	Frascati	?	S	
	Gran Sasso	?	S	
	Torino	?	S	
	Heidelberg	?	S	
	Munich	?	S	
	Marburg	?	S	
	Karlsruhe	?	S	
LANL	KEK	?	S,T	J. McGill
TUNL	Heidelberg	>89	S,T	G. Mitchell
	Münster	>889	S,T	
U. of NMex	Karlsruhe	>90	T,S	D. Wolfe
	CERN	>90	T,S	
	Rutherford	>90	T,S	

INFORMATION OBTAINED
FROM DOE OFFICE

From	To	Date	Type	Contact
LBL	CERN(NA36)	>88	S,T	L. Schroeder
	CERN(WA80)	>88	S,T	A. Poskanzer
	CERN(NA34)	>88	S,T	C. Gruhn
ORNL	CERN(WA80)	>88	S,T	F. Plasil
BNL	CERN(NA34)	>88	S,T	S. Lindenbaum
Yale	GSI	>88	S,T	Greenberger
UCLA	?	?	?	G. Igo
LANL	?	?	?	J. Moss

United States Institutions that use the LAMPF at the Los Alamos National Laboratory

ABILENE CHRISTIAN UNIVERSITY
ACU STATION
ABILENE, TX 79699

ARGONNE NATIONAL LABORATORY
9700 S. CASS AVENUE
ARGONNE, ILLINOIS 60439

ARIZONA STATE UNIVERSITY
PHYSICS DEPARTMENT
TEMPE,AZ

ASSOCIATED WESTERN UNIVERSITIES, INC
142 EAST 200 SOUTH
SALT LAKE CITY, UTAH 84111

BALL STATE UNIVERSITY
MUNCIE, IN 47306

BATES LINEAR ACCELERATOR
MIDDLETON, MA 01949

BATTELLE PACIFIC NORTHWEST LAB.
P. O. BOX 999
RICHLAND, WASHINGTON 99352

BELL LABORATORIES
MURRAY HILL, NJ 07974

BOSTON UNIVERSITY
111 CUMMINGTON STREET
BOSTON, MA 02148

BRIGHAM YOUNG UNIVERSITY
PROVO, UTAH 84602

BROOKHAVEN NATIONAL LABORATORY
UPTON, NY 11973

BROOKLYN COLLEGE
BROOKLYN, NEW YORK 11210

CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CA 91125

CALIFORNIA POLYTECHNIC STATE UNIVERSITY
SAN LUIS OBISPO, CA 93407

CALIFORNIA STATE UNIVERSITY
NORTHRIDGE, CA 91330

CARNEGIE-MELLON UNIVERSITY
PITTSBURGH, PA 15213

CASE WESTERN RESERVE UNIVERSITY
CLEVELAND, OHIO 44106

CATHOLIC UNIVERSITY OF AMERICA
WASHINGTON, D.C. 20064

CENTRAL WASHINGTON UNIVERSITY
ELLENSBURG, WA 98926

CLARK UNIVERSITY
950 MAIN STREET
WORCESTER, MA 01610

CLARKSON UNIVERSITY
POTSDAM, NY 13676

COLLEGE OF WILLIAM & MARY
PHYSICS DEPARTMENT
WILLIAMSBURG,VA

COLLEGE OF WILLIAM AND MARY
PHYSICS DEPARTMENT
WILLIAMSBURG, VA 23185

COLORADO COLLEGE
COLORADO SPRINGS, COLORADO 80903

CONTINUOUS ELECTRON BEAM ACCELERATOR
FACILITY -CE
12070 JEFFERSON AVENUE
NEWPORT NEWS, VIRGINIA 23606

DREXEL UNIVERSITY
32ND/CHESTNUT STREETS
PHILADELPHIA, PA 19104

DUKE UNIVERSITY
DURHAM, NORTH CAROLINA 27706

ELMHURST COLLEGE
ELMHURST, IL 60126

FERMI NATIONAL ACCELERATOR LABORATORY
P. O. BOX 500
BATAVIA, IL 60510

FLORIDA INTERNATIONAL UNIVERSITY
MIAMI, FL 33199

FLORIDA STATE UNIVERSITY
TALLAHASSEE, FLORIDA 32306

FOREIGN SCIENCE/TECHNOLOGY CTR.
220 7TH ST. NE
CHARLOTTESVILLE, VA 22901

FRANKLIN/MARSHALL COLLEGE
BOX 3003
LANCASTER, PA 17604

GEORGE MASON UNIVERSITY
FAIRFAX, VA 22030

GEORGE WASHINGTON UNIVERSITY
725 21ST ST
WASHINGTON, DC

GEORGIA INSTITUTE OF TECHNOLOGY
ATLANTA, GA 30332

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS 02138

IBM WATSON RESEARCH CENTER
P. O. BOX 218
YORKTOWN HEIGHTS, NY 10598

IDAHO NATIONAL ENGINEERING LAB.
550 SECOND STREET
IDAHO FALLS, ID 83401

ILLINOIS INSTITUTE OF TECHNOLOGY
CHICAGO, IL 60616

ILLINOIS STATE UNIVERSITY
NORMAL, IL 61761

INDIANA UNIVERSITY
PHYSICS DEPARTMENT
BLOOMINGTON, IN 47405

IOWA STATE UNIVERSITY
AMES, IA 50011

JOHNS-HOPKINS UNIVERSITY
BALTIMORE, MARYLAND 21218

KANSAS STATE UNIVERSITY
MANHATTAN, KANSAS 66506

KENT STATE UNIVERSITY
KENT, OHIO 44242

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720

LAWRENCE LIVERMORE NATIONAL LABORATORY
P.O. BOX 808
LIVERMORE, CALIFORNIA 94550

LEHIGH UNIVERSITY
BETHLEHEM, PA 18015

LOS ALAMOS NATIONAL LABORATORY
MS H841
PO BOX 1663
LOS ALAMOS, NM

LOUISIANA STATE UNIVERSITY
BATON ROUGE, LA 70803

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE
CAMBRIDGE, MA 02139

MICHIGAN STATE UNIVERSITY
EAST LANSING, MI 48824

NASA/LEWIS RESEARCH CENTER
CLEVELAND, OHIO 44135

NATIONAL BUREAU OF STANDARDS
WASHINGTON, D.C. 20234

NATIONAL SCIENCE FOUNDATION
1800 G STREET, NW
WASHINGTON, D. C. 20550

NEW MEXICO INSTITUTE OF MINING/TECHNOLOGY
CAMPUS STATION
SOCORRO, NM 87801

NEW MEXICO STATE UNIVERSITY
PHYSICS DEPARTMENT
LAS CRUCES, N M 8803

NEW YORK UNIVERSITY
UNIVERSITY HEIGHTS
THE BRONX, NEW YORK 10453

NORTHWESTERN UNIVERSITY
EVANSTON, ILLINOIS 60210

OAK RIDGE NATIONAL LABORATORY
P.O. BOX
OAK RIDGE, TENNESSEE 37830

OHIO STATE UNIVERSITY
PHYSICS DEPARTMENT
174 W 18TH ST
COLUMBUS, OHIO 43210

OHIO UNIVERSITY
ATHENS, OHIO 45701

OREGON STATE UNIVERSITY
CORVALLIS, OR 97331

PENNSYLVANIA STATE UNIVERSITY
UNIVERSITY PARK, PA 16802

PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08540

PURDUE UNIVERSITY
PHYSICS DEPARTMENT
WEST LAFAYETTE, IN 47907

RENSSELAER POLYTECHNIC INSTITUTE
TROY, NEW YORK 12181

RICE UNIVERSITY
HOUSTON, TX 77251

ROCKWELL HANFORD OPERATIONS
202-S 200-W
P.O. BOX 800
RICHLAND, WA 99352

RUTGERS UNIVERSITY
PHYSICS DEPT
PISCATAWAY, NJ

SAN JOSE STATE UNIVERSITY
SAN JOSE, CA 95192

SANDIA LABORATORIES, ALBQ
ALBUQUERQUE, NEW MEXICO 87185

SANDIA LABORATORIES, LIVERMORE
LIVERMORE, CA 94550

SCRIPPS INSTITUTION OF OCEANOGRAPHY
LA JOLLA, CA 92093

SOUTHEASTERN UNIVERSITY RESEARCH ASSOCIATION
12070 JEFFERSON AVE.
NEWPORT NEWS, VA 23606

SOUTHWEST RESEARCH INSTITUTE
P.O. DRAWER 28510
SAN ANTONIO, TX 78284

STANFORD LINEAR ACCELERATOR CENTER
P.O. BOX 4349
STANFORD, CALIFORNIA 94305

STANFORD UNIVERSITY
PHYSICS DEPARTMENT
VARIAN PHYSICS BLDG
STANFORD, CA 94305

STATE UNIVERSITY OF NEW YORK, ALBANY
1400 WASHINGTON AVENUE
ALBANY, NY 12222

SYRACUSE UNIVERSITY
SYRACUSE, NY 13210

TEMPLE UNIVERSITY
PHILADELPHIA, PA 19122

TEXAS A/M UNIVERSITY
COLLEGE STATION, TEXAS 77843

TEXAS TECH UNIVERSITY
LUBBOCK, TX 79409

U.S. AIR FORCE
ALBUQUERQUE, NM

U.S. DEPARTMENT OF ENERGY
WASHINGTON, DC 20545

U.S. NAVAL ACADEMY
ANNAPOLIS, MD 21402

UNIVERSITY OF ALABAMA
P.O.BOX 1921
UNIVERSITY, AL 35486

UNIVERSITY OF ALBERTA
EDMONTON, ALBERTA
CANADA T6G 2J1

UNIVERSITY OF ARIZONA
TUCSON, AZ 85721

UNIVERSITY OF ARKANSAS
PHYSICS DEPARTMENT
PINE BLUFF, AK

UNIVERSITY OF CALIFORNIA AT LOS ANGELES
PHYSICS DEPT
LOS ANGELES, CA

UNIVERSITY OF CALIFORNIA, IRVINE
IRVINE, CA 92717

UNIVERSITY OF CALIFORNIA, LOS ANGELES
LOS ANGELES, CA 90024

UNIVERSITY OF CALIFORNIA, RIVERSIDE
RIVERSIDE, CA 92502

UNIVERSITY OF CALIFORNIA, SAN DIEGO
P.O. BOX 109
LA JOLLA, CA 92093

UNIVERSITY OF CALIFORNIA, SANTA BARBARA
SANTA BARBARA, CA 93106

UNIVERSITY OF CALIFORNIA, SANTA CRUZ
SANTA CRUZ, CA 95064

UNIVERSITY OF CHICAGO
5630 S. ELLIS
CHICAGO, IL 60637

UNIVERSITY OF CINCINNATI
CINCINNATI, OH 45221

UNIVERSITY OF COLORADO
BOULDER, CO 80309

UNIVERSITY OF CONNECTICUT
STORRS, CT 06268

UNIVERSITY OF DENVER
DENVER, CO 80210

UNIVERSITY OF GEORGIA
ATHENS, GA 30602

UNIVERSITY OF HOUSTON
HOUSTON, TX 77004

UNIVERSITY OF IDAHO
MOSCOW, IDAHO 83843

UNIVERSITY OF ILLINOIS
URBANA, ILLINOIS 61801

UNIVERSITY OF IOWA
IOWA CITY, IA 52242

UNIVERSITY OF KANSAS
LAWRENCE, KANSAS 66044

UNIVERSITY OF KENTUCKY
LEXINGTON, KY 40506

UNIVERSITY OF LOUISVILLE
LOUISVILLE, KY 40292

UNIVERSITY OF MARYLAND
COLLEGE PARK, MD 20742

UNIVERSITY OF MASSACHUSETTS
AMHERST, MASSACHUSETTS 01003

UNIVERSITY OF MICHIGAN
PHYSICS DEPARTMENT
ANN ARBOR, MI 48109

MINNEAPOLIS, MINNESOTA
PHYSICS DEPARTMENT
116 CHURCH ST
MINNEAPOLIS, MN

UNIVERSITY OF MISSISSIPPI
UNIVERSITY, MS 38677

UNIVERSITY OF MISSOURI
COLUMBIA, MO 65211

UNIVERSITY OF MONTANA
MISSOULA, MONTANA 59812

UNIVERSITY OF NEW HAMPSHIRE
DURHAM, NH 03824

UNIVERSITY OF NEW MEXICO
ALBUQUERQUE, NM 87131

UNIVERSITY OF OKLAHOMA
PHYSICS DEPARTMENT
440 W. BROOKS
NORMAN, OKLAHOMA 73019

UNIVERSITY OF OREGON
EUGENE, OREGON 97403

UNIVERSITY OF PENNSYLVANIA
PHYSICS DEPARTMENT
209 S. 33RD ST

PHILADELPHIA, PA 19104

UNIVERSITY OF PHILADELPHIA
209 S. 33RD ST.
PHILADELPHIA, PA

UNIVERSITY OF PITTSBURGH
PITTSBURGH, PA 15260

UNIVERSITY OF ROCHESTER
ROCHESTER, NY 14627

UNIVERSITY OF SOUTH CAROLINA
COLUMBIA, SOUTH CAROLINA 29208

UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90007

UNIVERSITY OF TEXAS
PHYSICS DEPARTMENT
AUSTIN, TX 78712

UNIVERSITY OF TEXAS, ARLINGTON
P.O. BOX 19363
ARLINGTON, TX 76014

UNIVERSITY OF UTAH
SALT LAKE CITY, UTAH

UNIVERSITY OF VIRGINIA
CHARLOTTESVILLE, VA 22901

UNIVERSITY OF WASHINGTON
SEATTLE, WA 98195

UNIVERSITY OF WYOMING
LARAMIE, WYOMING 82071

UTAH STATE UNIVERSITY
PHYSICS DEPT
LOGAN,UT 84322

VALPARAISO UNIVERSITY
PHYSICS DEPT
VALPRAISO, IN 46383

VASSAR COLLEGE
POUGHKEEPSIE, NY 12601

VIRGINIA STATE UNIVERSITY
PETERSBURG, VA 23803

VPI/STATE UNIVERSITY
BLACKSBURG, VA 24061

WASHINGTON STATE UNIVERSITY
PULLMAN, WA 99164

WORCESTER POLYTECHNIC INSTITUTE
WORCESTER, MA 01609

YALE UNIVERSITY
217 PROSPECT AVE
NEW HAVEN, CT

Foreign Institutes that use the LAMPF facility at the Los Alamos National Laboratory.

BEN-GURION UNIVERSITY
BEERSHEVA, ISRAEL 84110

** C.E.N. SACLAY
91190 GIF-SUR-YVETTE
FRANCE

C.N.R.S.
15 QUAI ANATOLE FRANCE
PARIS, FRANCE

CENTRAL RESEARCH LAB. OF HITACHI,LTD.
TOKYO 185, JAPAN

CENTRE DE PHYSIQUE DES PARTICULES
FACULTE DES SCIENCES DE LUMINY
CASE 907 13288 MARSEILLE
CEDEX 9, FRANCE

CENTRE NATIONAL/RECHERCHE SCIENTIFIQUE
15 QUAI ANATOLE FRANCE
75007-PARIS
FRANCE

** CERN
1211 GENEVE 23
SWITZERLAND

CHALK RIVER NUCLEAR LABORATORIES
CHALK RIVER
ONTARIO, CANADA K0J 1J0

CHALMERS UNIVERSITY OF TECHNOLOGY
S-412 96 GOTEBORG
SWEDEN

CHANGSHA INSTITUTE OF TECHNOLOGY
NO. 137,YAN WA CHI STREET
CHANGSHA, HUNAN PROVINCE
PEOPLES REPUBLIC OF CHINA

CHUKYO COLLEGE
TOKI-CHO, MIZUNAMI-SHI
GITU-KEN, JAPAN

** DESY
NOTKESTR. 85
2000 HAMBURG 52
WEST GERMANY

EIDGENOSSISCHE TECHNISCHE
HOCHSCHULE ZURICH
HONGGERBERG, CH-8093
ZURICH, SWITZERLAND

EINDHOVEN UNIVERSITY OF TECHNOLOGY
DEN-DOLECH 2
EINDHOVEN 5600MB
THE NETHERLANDS

ETH ZURICH
HONGGERBERG
CH-8093 ZURICH, SWITZERLAND

FIRST NAZIONALE FISICA NUCLEARE
SEZIONE DE CATANIA
1 95129 CATANIA
CORSO, ITALY 57

FLINDERS UNIVERSITY OF SOUTH AUSTRALIA
BEDFORD PARK, SA 5042, AUSTRALIA

** GSI
POSTFACH 110541
D-6100 DARMSTADT
WEST GERMANY

** GUSTAF WERNER INSTITUTE
S-75121 UPPSALA
SWEDEN

HACHINOHE INSTITUTE OF TECHNOLOGY
HACHINOHE, AOMORI
031 JAPAN

HAHN-MEITNER-INSTITUT
GLIENICKERSTRASSE 100
1 BERLIN 39
WEST GERMANY

HEBREW UNIVERSITY
RACAH INST OF PHYSICS
JERUSALEM 91904
ISRAEL

NATIONAL LABORATORY FOR HIGH ENERGY PHYSICS
COUNTER EXPERIMENT DIVISION
OHO-MACHI, TSUKUBA-GUN
IBARAKI-KEN, 300-32, JAPAN

** NATIONAL RESEARCH COUNCIL OF CANADA
OTTAWA, ONTARIO
CANADA K1A 0R6

NIHON UNIVERSITY
KANDA-SURUGADAI, CHIYODA-KU
TOKYO, JAPAN

NIKHEF
P. O. BOX 4395
1109 AJ AMSTERDAM
THE NETHERLANDS

PHILIPPS UNIVERSITAT, LAHNBERGE
D-355 MARBURG
WEST GERMANY

RISO NATIONAL LABORATORY
P. O. BOX 49
DK-4000 ROSKILDE
DENMARK

ROGALAND RESEARCH INSTITUTE
P. O. BOX 2503
ULLANDHANG, N-4001
NORWAY

ROYAL INSTITUTE OF TECHNOLOGY, SWEDEN
10044, STOCKHOLM, SWEDEN

RUDER BOSKOVIC INSTITUTE
PO BOX 1016
41001 ZAGREB, BIJENICKA 54
CROATIA,YUGOSLAVIA

** RUTHERFORD APPLETON LABORATORY
CHILTON, DIDCOT
OXON, U.K. OX11 0QX

SAITAMA UNIVERSITY
URAWA CITY
SAITAMA-KEN, JAPAN 336

SHEFFIELD UNIVERSITY
WESTERN BANK
SHEFFIELD, ENGLAND

** SIMON FRASER UNIVERSITY
BURNABY, B.C.
CANADA V5A 1S6

** SIN
CH-5234 VILLIGEN
ZURICH, SWITZERLAND

TATA INSTITUTE OF FUNDAMENTAL RESEARCH
BOMBAY, INDIA

TECHNISCHE UNIVERSITAT MUNCHEN
PHYSIK DEPARTMENT, E 18
D8046 GARCHING
WEST GERMANY

TEIKYO UNIV. SCHOOL OF MEDICINE
2-11-1, KAGA, ITABASHI-KU
TOKYO 173, JAPAN

** TEL-AVIV UNIVERSITY
RAMAT AVIV, ISRAEL 69978

TOHOKU UNIVERSITY
AZA-AOBA
SENDAI, JAPAN

TOKYO INSTITUTE OF TECHNOLOGY
OH-OKAYAMA, NEGURO
TOKYO, JAPAN

TOKYO METROPOLITAN UNIVERSITY
2-1-1 FUKAZAWA, SETAGAYA-KU
TOKYO 158, JAPAN

UNIVERSIDAD SIMON BOLIVAR
CARACAS, VENEZUELA

UNIVERSITA DEGLI STUDI
VIA LARGA
33100 UDINE, ITALY

UNIVERSITA DEGLI STUDY
INSTITUTE DE FISICA
VIA LARGA 36
33100 UDINE, ITALY

** UNIVERSITAT HEIDELBERG
D-6900 HEIDELBERG-1
WEST GERMANY

HIROSHIMA UNIVERSITY
HIROSHIMA 730, JAPAN

HOCHSCHULE DER BUNDESWEHR MUNCHEN
WERNER-HEISENBERG-WEG 39
8014 NEUBIBERG
WEST GERMANY

INST. FUR RADIUMFORSCHUNG AND KERNPHYSIK
BOLTZMANNG 3
A-1040 VIENNA, AUSTRIA

INST. NAZIONALE DI FISICA NUCLEARE
VIA CELORIA 16
I20133 MILANO, ITALY

INSTITUT DE PHYSIQUE NUCLEAIRE
B.P. NO. 1
91406 ORSAY, CEDEX, FRANCE

** INSTITUT FUR KERNPHYSIK
JULICH
WEST GERMANY

INSTITUTE OF ATOMIC ENERGY, PRC
ACADEMIA SINICA
PEKING
PEOPLES REPUBLIC OF CHINA

INSTITUTE OF HIGH ENERGY PHYSICS, PRC
ACADEMIA SENECA
PO BOX 918, PEKING
PEOPLE'S REPUBLIC OF CHINA

INSTITUTE OF NUCLEAR RESEARCH, USSR
ACADEMY OF SCIENCES
PROFSOJUZNAJA 7A
MOSCOW, USSR

INSTITUTE OF THEORETICAL PHYSICS, PRC
ACADEMIA SINICA
P. O. BOX 2735 BEIJING
PEOPLE'S REPUBLIC OF CHINA

INSTITUTE R. BOSKOVIC
P. O. BOX 1016
41001 ZAGREB, YUGOSLAVIA

ISRAEL ATOMIC ENERGY COMMISSION
SOREQ NUCL. RESEARCH CENTRE
YAVNE, ISRAEL

JET
ABINGDON, OX14 3EA
ENGLAND

JORDAN UNIVERSITY
AMMAN, JORDAN

** KEK
NATL LAB FOR HIGH ENERGY PHYSICS
OHO-MACHI,TSUKUBA-GUN,
IBARAKI-KEN, 305 JAPAN

KERNFYSISCH VERSNELLER INSTITUUT
UNIVERSITEITS COMPLEX PADDEPOEL
GRONINGEN, THE NETHERLANDS

KFA
POSTFACH 1913
5170 JULICH 1
WEST GERMANY

KYOTO SANGYO UNIVERSITY
KAMIGAMO-MOTOYAMA
KYOTO 603, JAPAN

** KYOTO UNIVERSITY
UJI, KYOTO 611, JAPAN

LABORATORIO DE ESTUDOS AVANCADOS
SAO JOSE DOS CAMPOS
SAO PAULO, BRAZIL 12200

LABORATORIUM FUR KERNPHYSIK
ETH-HONGGERBERG
CH-8093 ZURICH
SWITZERLAND

LAKEHEAD UNIVERSITY
THUNDER BAY
ONTARIO, CANADA P7B 5E1

MAX-PLANCK INSTITUT FUR KERNPHYSIK
POSTFACH 103 980
D-6900 HEIDELBERG 1
WEST GERMANY

MAX-PLANCK-INSTITUTE
D-8000 MUNICH 23
WEST GERMANY

NAGOYA UNIVERSITY
FURO-CHO, CHIKUSA-KU
NAGOYA 464 JAPAN

UNIVERSITAT MUNCHEN
AM COULOMBWASSE
8046 GARCHING
WEST GERMANY

UNIVERSITE CATHOLIQUE DE LOUVAIN
INST PHYSIQUE/CHEMIN DU CYCL 2
B-1348 LOUVAIN-LA-NEUVE
BELGIUM

** UNIVERSITE DE GRENOBLE
53,AVENUE DES MARTYRS
38026 GRENOBLES, FRANCE

UNIVERSITE DE NEUCHATEL
2000 NEUCHATEL, SWITZERLAND

UNIVERSITE LAVAL - PHISIQUE
CITE UNIVERSITAIRE
QUEBEC, CANADA G1K 7P4

UNIVERSITET 1 TRONDHEIM
FYSISK INSTITUTT
NORGES LAERERHOGSKOLE
7000 TRONDHEIM, NORWAY

UNIVERSITY DE CLERMONT-FERRAND
B.P. NO. 45
AUBIERE, FRANCE

UNIVERSITY GIESSEN
2 PHYSIK INSTITUT
ARNDSTR. 2, D-63 GIESSEN
WEST GERMANY

UNIVERSITY MAINZ
INSTITUTE F. KERNPHYSIK
D65 MAINZ
WEST GERMANY

UNIVERSITY OF ADELAIDE
P. O. BOX 498
G.P.O. ADELAIDE SOUTH AUSTRALIA
AUSTRALIA 5001

UNIVERSITY OF BIRMINGHAM
P.O. BOX 363
BIRMINGHAM B15 2TT
ENGLAND

UNIVERSITY OF BERNE
PHYSICS INSTITUTE
SIDLERSTR. 5
CH-3012 BERN, SWITZERLAND

UNIVERSITY OF BONN
D-5300 BONN
WEST GERMANY

** UNIVERSITY OF BRITISH COLUMBIA, TRIUMF
VANCOUVER, B.C.
CANADA V6T 2A3

UNIVERSITY OF ERLANGEN
INSTITUTE OF THEORETICAL PHYSICS
8520 ERLANGEN
WEST GERMANY

UNIVERSITY OF FREIBURG, FAK. F. PHYSIKS
HERMANN-HERDER-STRASSE 3
FREIBURG, I. BR.
WEST GERMANY

UNIVERSITY OF FRIBOURG
CH-1700 FRIBOURG, SWITZERLAND

** UNIVERSITY OF GENEVA
1211 GENEVE 4
SWITZERLAND

UNIVERSITY OF GLASGOW
GLASGOW G12 8QQ, SCOTLAND

UNIVERSITY OF HEIDELBERG
PHILOSOPHENWEG 12
69 HEIDELBERG
WEST GERMANY

UNIVERSITY OF INNSBRUCK
6020 INNSBRUCK
AUSTRIA

UNIVERSITY OF KARLSRUHE
POSTFACH 3640
75 KARLSHRUHE, WEST GERMANY

UNIVERSITY OF LONDON
REGENTS PARK
LONDON NW1 4NS, ENGLAND

** UNIVERSITY OF MANITOBA
WINNIPEG, MANITOBA
CANADA R3T 2N2

UNIVERSITY OF MARBURG
MAINZERGASSE 33
3550 MARBURG
WEST GERMANY

UNIVERSITY OF MELBOURNE
PARKSVILLE, VICTORIA
AUSTRALIA 3104

UNIVERSITY OF PARIS
24 RUUE LHOMOND
75231 PARIS CEDEX 05, FRANCE

UNIVERSITY OF PETROLEUM/MINERALS
DHARAN, SAUDI ARABIA 31261

UNIVERSITY OF REGINA
REGINA, SASKATCHEWAN
CANADA

UNIVERSITY OF SIEGEN
59 SIEGEN 21
WEST GERMANY

UNIVERSITY OF SURREY
GUILDFORD, SURREY
UNITED KINGDOM

UNIVERSITY OF THE WITWATERSRAND
JOHANNESBURG, SOUTH AFRICA

UNIVERSITY OF TOKYO
TOKYO, JAPAN

UNIVERSITY OF TORINO
DELL UNIVERSITA DI TORINO
TORINO, ITALY 10125

UNIVERSITY OF TUBINGEN
TUBINGEN, WEST GERMANY

UNIVERSITY OF TURKU
SF-20500
TURKU 50, FINLAND

UNIVERSITY OF UDINE
VIA MANTICA 3
I-33100 UDINE
ITALY

UNIVERSITY OF YORK
HESLINGTON YORK
YO1 5DD
ENGLAND

UNIVERSITY OF UTRECHT
PRINCETONPLEIN 5
P.B. 80.006
THE NETHERLANDS

UNIVERSITY OF VICTORIA
VICTORIA, B.C., CANADA V8W 2Y2

UNIVERSITY OF VIENNA
A-1090, BOLTZMANNG. 3
VIENNA, AUSTRIA

UNIVERSITY OF WEST ONTARIO
LONDON, ONTARIO
CANADA N6A 3K7

UNIVERSITY OF ZAGREB
P.O. BOX 162
41001 ZAGREB, YOGOSLAVIA

** UNIVERSITY OF ZURICH
SCHOENBERGGASSE 9
CH-8001 ZURICH, SWITZERLAND

VRIJE UNIVERSITY
POSTBUS 7161 MC
AMSTERDAM
THE NETHERLANDS

**WEIZMANN INSTITUTE OF SCIENCE
76100 REHOVOT
ISRAEL

FUSION ENERGY

NETWORK REQUIREMENTS

MFEnet II
FIVE YEAR NETWORKING PLAN
PHILOSOPHY

J.F.Leighton
6 JAN 1987

CONTENTS

ABSTRACT: A Philosophy Model

This document describes an outline or framework upon which the five year planning for the MFEnet II is based. It is intended to present more of the philosophy of the planning, emphasing the overall structure of the plan, rather than the details and time frame of such a plan.

INTRODUCTION

The National Magnetic Fusion Energy Computer Center (NMFECC) is located at the Lawrence Livermore National Laboratory near Livermore, California. The purpose of this Computing Center is to provide the resources needed to support the large scale computing needs of the Office of Energy Research in the Department of Energy. The Center maintains and operates a system of supercomputers, (currently two CRAY 1s, a CRAY XMP-22, and a CRAY 2) to supply the major computing resources. These computing resources are made available to researchers throughout the United States and at selected locations abroad though a data communications network called the MFEnet. (See Figures I,II)

A new effort in data communications networking, designated ESnet (Energy Sciences Network) has been initiated by the Office of Energy Research (OER) under the mangement of the OER Scientific Computing Staff. The ESnet concept is intended to bring together the various data communications efforts of the different research programs funded by the OER into a single coherent effort. The goal is to improve the networking capabilities of these programs, and at the same time, to reduce costs associated with redundant or incompatible data communcations projects.

The NMFECC networking staff has proposed, and been funded for, a new network design (currently called MFEnet II) as a major component of the ESnet project. MFEnet II will be designed , implemented, and operated by the NMFECC network staff. This new design will gradually replace the current MFEnet. It will require an estimated three years to complete the phaseover of the current user community of MFEnet to

MFEnet II. The additional requirements of the ESnet effort must also be addressed during that timeframe.

AN OVERVIEW OF MFEnet II

A conceptual diagram of the MFEnet II architecture is shown in Figure III. This architecture is intended to met a number of requirements including presenting to the user a more "open" architecture. The network will connect to a user site's ethernet local area network (LAN) and will support any higher level protocols which use the ARPA Internet Protocol, IP. A second means of interfacing will be via (directly connected) X.25 ports. Both IP and X.25 are widespread standard (or pseudo-standard protocols) and commercially available. Additionally, the ISO Internet Protocol (also called IP) will be supported at a later time. The initial IP capability will also facilitate connections to other IP based networks, the ARPA Internet being one of the more prominent.

The NSP transport level protocol used by the NMFECC CRAYs will be modified to use IP as part of the implementation of MFEnet II. The project also requires the provision of several types of NSP/IP based server which will be able to connect to the user's LAN. These servers are intended to provide access to the services of the (NSP/IP based) CRAY supercomputers.

NETWORK PLANNING CONSIDERATIONS

In the development of a long range plan for the MFEnet II effort, it became evident that it would be instructional to momentarily go back to basic issues and try to identify the fundamentals of what a data communications network should be supplying to it's user community. A somewhat high level, abstract, three-level network support model has resulted from this thought process. The three levels of support identified are:

<u>Connectivity</u>: Providing the physical connections to the Network for the user community as well as the corresponding network infrastructure.

Interoperability: Providing the ability for separate protocol user groups to communicate with each other.

This may be done in several ways:

1. Gateways
2. Use of Multiple protocol suites within a community
3. Convergence to a common protocol suite

Functionality: Providing new capabilities for communications, resource sharing, and remote or distributed computing.

One significance of these three levels of support to our planning is that although all three must be addressed and planned for, all three are not solely "network" issues.

CORRELATION TO THE ISO MODEL

These three areas of consideration can be mapped onto the International Standards Organization (ISO) Open Systems Interconnect (OSI) model. This is shown in Figure IV. The figure shows the seven protocol layers of the OSI model, the major issue dealt with at each layer, the location of each layer (in the MFEnet II implementation), examples of each layer, and finally how our three planning levels of support correlate with the layers of the ISO model.

The intent of this diagram is to help clarify the areas of concern when planning for providing the above three levels of support. Since it is obvious that each level is supported by the level beneath it, our plan will be to work from the bottom up - concentrating first on connectivity, then interoperability, and finally addressing functionality. It should be noted that there are short term, temporary "fixes" in each area, and therefore various degrees of support will be present at all three levels throughout the planning period.

PLANNING: SHORT RANGE (1-2 years)

Short range planning will be directed at providing full connectivity. This will be done principally through the vehicle of implementing and expanding MFEnet II. The user community on the current MFEnet will be gradually moved over to MFEnet II, new international links will be installed, and selected links will be upgraded in performance.

- CONNECTIVITY
  - Closed to Open Network Access:
    - Ethernet connection
    - IP and/or X.25 Access
  - Phaseover of MFE Community
  - Interconnection of new communities
  - International Connections
    - Japan/W.Germany/Switzerland
  - Performance Enhancement of some links

Interoperability will be achieved by using the most readily available means: using multiple protocol suites in a given host and, for electronic mail, using existing mail gateways.

- INTEROPERABILITY
  - Multiple protocol suites
  - Mail Gateways

Functionality provided will be essentially that which is currently available.

- FUNCTIONALITY
  - Standard Functions:
    - Remote Terminal Access
    - File Transfer
    - Electronic Mail
  - Enhanced Servers and Naming Conventions
  - Centralized Network Operations
    - Privacy,Access Control,Accounting,Statistics
  - MFECC and FSU Supercomputer Access

PLANNING: MEDIUM RANGE (2-4 years)

Medium term planning will continue with enhancing connectivity. The gradual phaseover of the MFEnet community to MFEnet II should be completed. Newer sites will continue to be added. The activities of implementing a governmental interagency internetwork will be started, as well as a restructuring of the MFEnet II communications infrastructure*. The MFEnet II gateways will be enhanced to be compatible with the ISO Internet Protocol (ISO IP). We will also be exploring means of extending into the home environment the network support that researchers have available in their office.

CONNECTIVITY
- Final Phaseover of MFE Community
- Interconnection of new communities
- Interagency Internetworking
- High Bandwidth Broadcast (1M bit/sec)
- CompleteTerrestrial interconnectivity
- ISO IP routing
- Extensions to home environment

Interoperability will receive additional attention. One effort will be to provide TCP/IP access to the NMFECC CRAYS in parallel with the NSP/IP protocol suite. This period will also mark the beginning of the transition to the ISO standard protocol set. However, much of the interoperability will still be achieved by maintaining multiple protocol sets in the host computers.

INTEROPERABILITY
- TCP/NSP application level gateways
  -or- TCP/IP on Crays
- Introduction of ISO protocols
- Multiple protocol suites

Enhanced functionality will be based largely on distributing services across the network. There will be some specialized distributed processing applications available. The framework for more general distributed applications will be underway in the form of remote procedure call support, possible an implementation of the industry standard Network File System, and a general purpose task-to-task communication capability. Some of the new capabilities of the network communications infrastructure will require type-of-service access at the higher protocol levels. We will continue to enhance the office environment through PC and workstation support.

FUNCTIONALITY
- Specialized Distributed Processing
- Remote Procedure Calls
- Network File System Support
- Task-to-Task Communication
- Type-of-Service Routing
- PC as host
- Workstation Support

---

[8] The MFEnet II communications infrastructure uses a mixture of terrestrial and satellite communications links. A modified structure which supplies full interconnectivity via high speed (but long delay) satellite links, and also full interconnectivity via terrestrial links is shown in Figure V. This approach offers both high bandwidth and low delay paths, and the ability to route appropriate traffic accordingly. Additionally, the two path sets offer back-up paths for each other in case of link outage.

PLANNING: LONG RANGE (3-5 years)

Long term connectivity will be a continuation of previous issues, with a on-going addition of new users, and upgrading on new links. Some projections for this time-frame are speculating on applications the will require much higher bandwidth than is presently used.

CONNECTIVITY
Very High Bandwidth Broadcast (above1M bit/sec)

The issues of interoperability should diminish in the long term as the transition to the ISO standard protocols in nearly complete. Even with the availability of ISO protocols, we intend to also support the NSP/IP protocol set as a vehicle for doing research in networking.

INTEROPERABILITY
Transition to ISO protocols

The functionality to be supported in the longer term may include some of the following applications.

FUNCTIONALITY
- ? Multimedia Mail
- ? Multimedia Conferencing
- ? Industry Standard Windowing
- ? Remote Instrument Support
- ? Remote Data Acquisition
- ? Generalized Distributed Processing

SUMMARY

This document has emphasized a philosophy upon which to base a five year plan, rather that the details of the plan itself. A long range plan should be viewed as a means of setting the direction for the future, as opposed to establishing the path. Hopefully, our philosophy for the long range will remain fairly stable, although the plan will undoubtedly change.

National MFE Network 1986

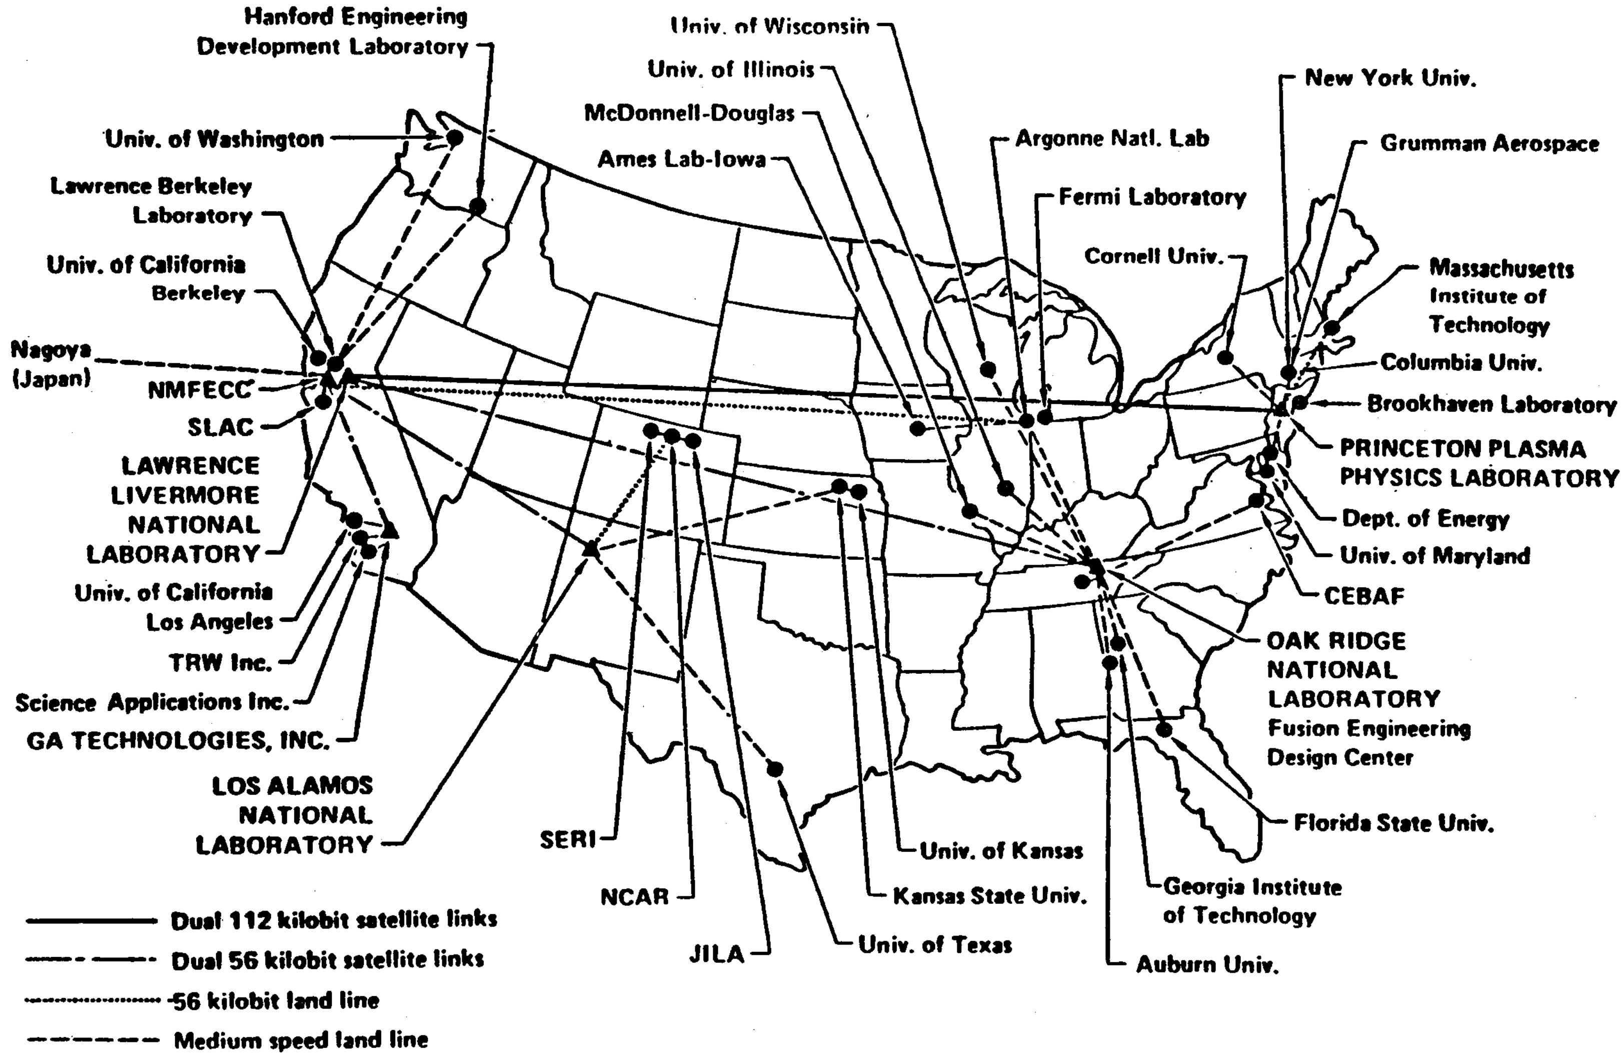

FIGURE I

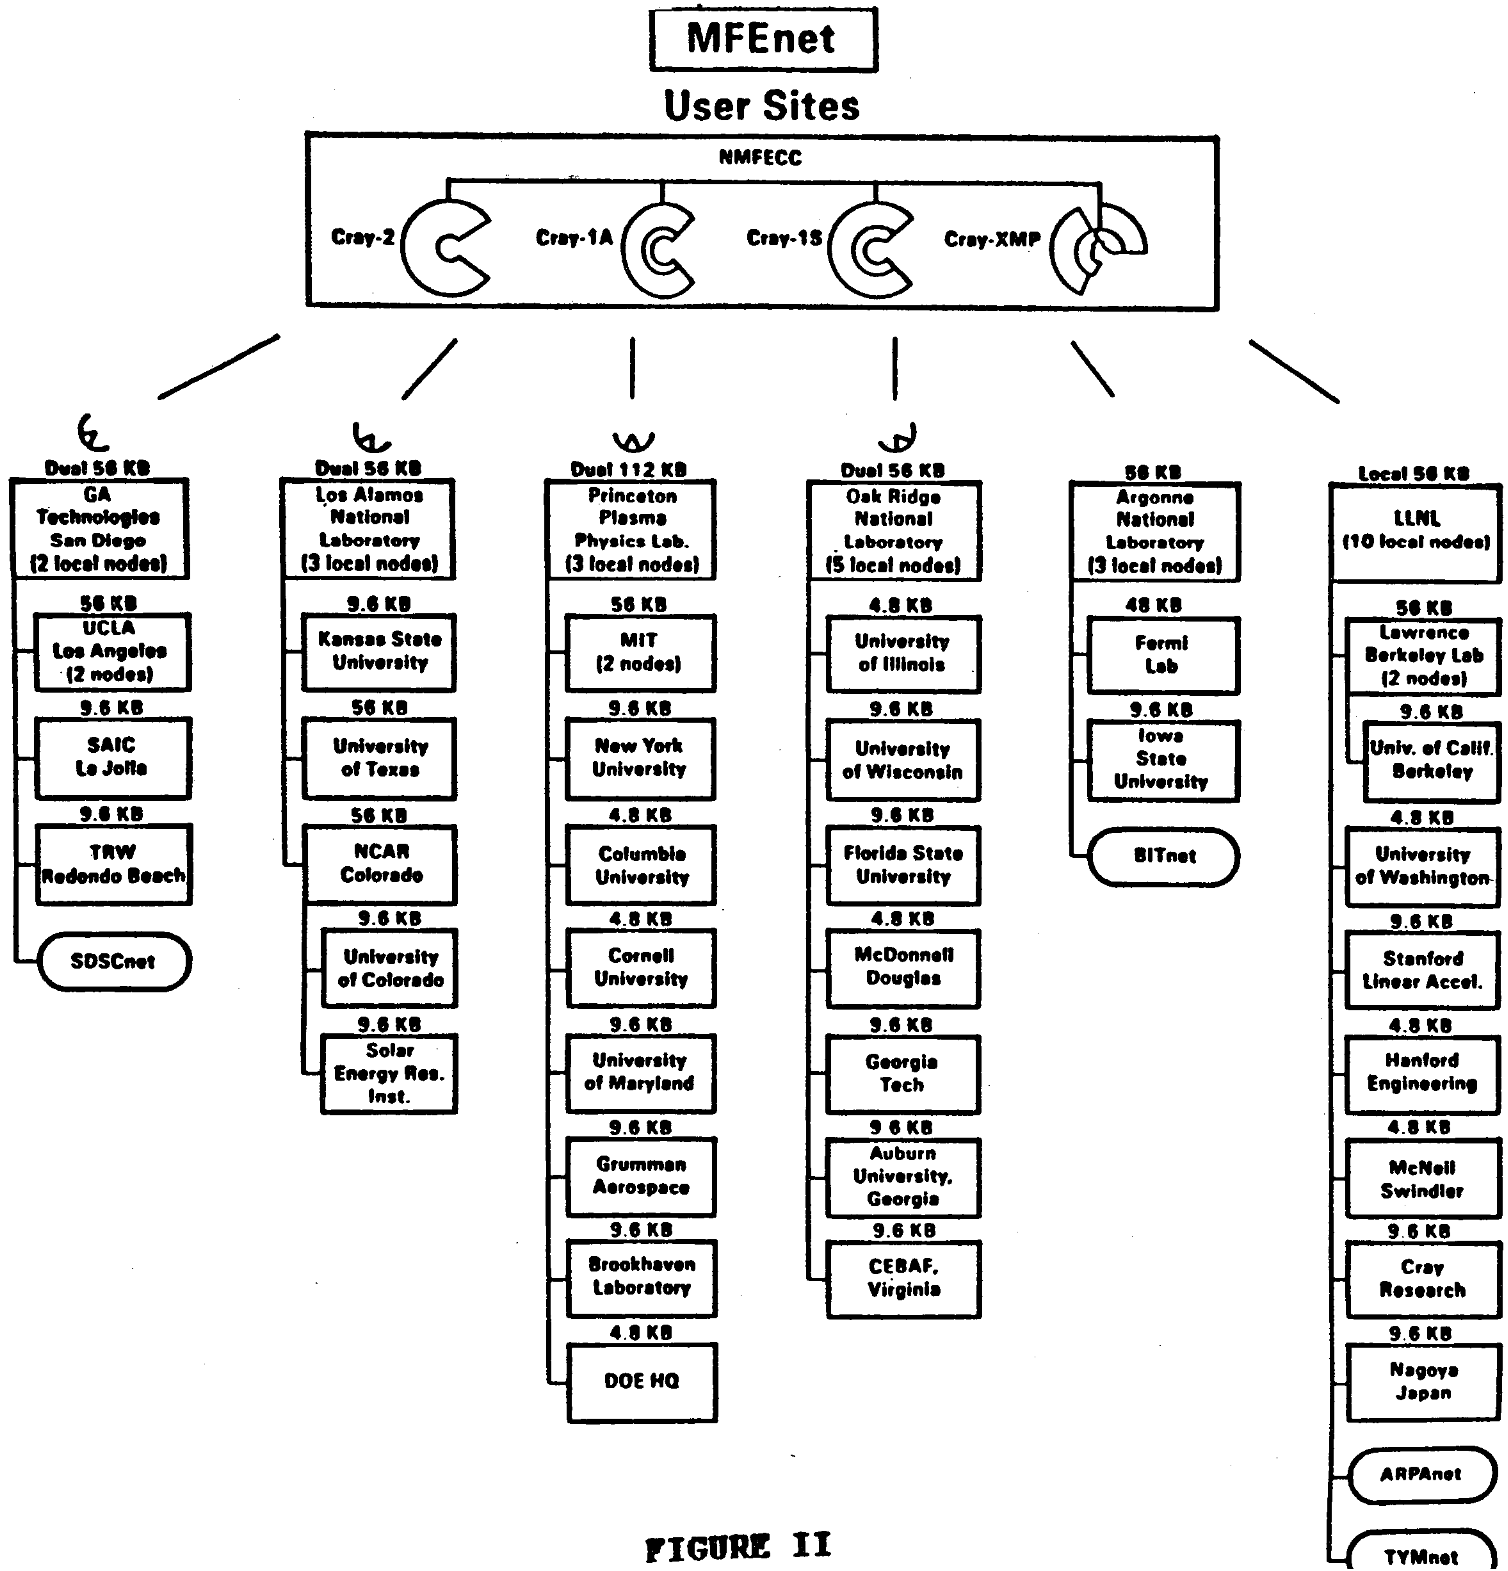
MFEnet
User Sites
NMFECC
Cray-2
Cray-1A
Cray-1S
Cray-XMP
Dual 56 KB
GA Technologies San Diego (2 local nodes)
56 KB
UCLA Los Angeles (2 nodes)
9.6 KB
SAIC La Jolla
9.6 KB
TRW Redondo Beach
SDSCnet
Dual 56 KB
Los Alamos National Laboratory (3 local nodes)
9.6 KB
Kansas State University
56 KB
University of Texas
56 KB
NCAR Colorado
9.6 KB
University of Colorado
9.6 KB
Solar Energy Res. Inst.
Dual 112 KB
Princeton Plasma Physics Lab. (3 local nodes)
56 KB
MIT (2 nodes)
9.6 KB
New York University
4.8 KB
Columbia University
4.8 KB
Cornell University
9.6 KB
University of Maryland
9.6 KB
Grumman Aerospace
9.6 KB
Brookhaven Laboratory
4.8 KB
DOE HQ
Dual 56 KB
Oak Ridge National Laboratory (5 local nodes)
4.8 KB
University of Illinois
9.6 KB
University of Wisconsin
9.6 KB
Florida State University
4.8 KB
McDonnell Douglas
9.6 KB
Georgia Tech
9.6 KB
Auburn University, Georgia
9.6 KB
CEBAF, Virginia
56 KB
Argonne National Laboratory (3 local nodes)
48 KB
Fermi Lab
9.6 KB
Iowa State University
BITnet
Local 56 KB
LLNL (10 local nodes)
56 KB
Lawrence Berkeley Lab (2 nodes)
9.6 KB
Univ. of Calif. Berkeley
4.8 KB
University of Washington
9.6 KB
Stanford Linear Accel.
4.8 KB
Hanford Engineering
4.8 KB
McNeil Swindler
9.6 KB
Cray Research
9.6 KB
Nagoya Japan
ARPAnet
TYMnet

FIGURE II

MFEnet II
Block diagram

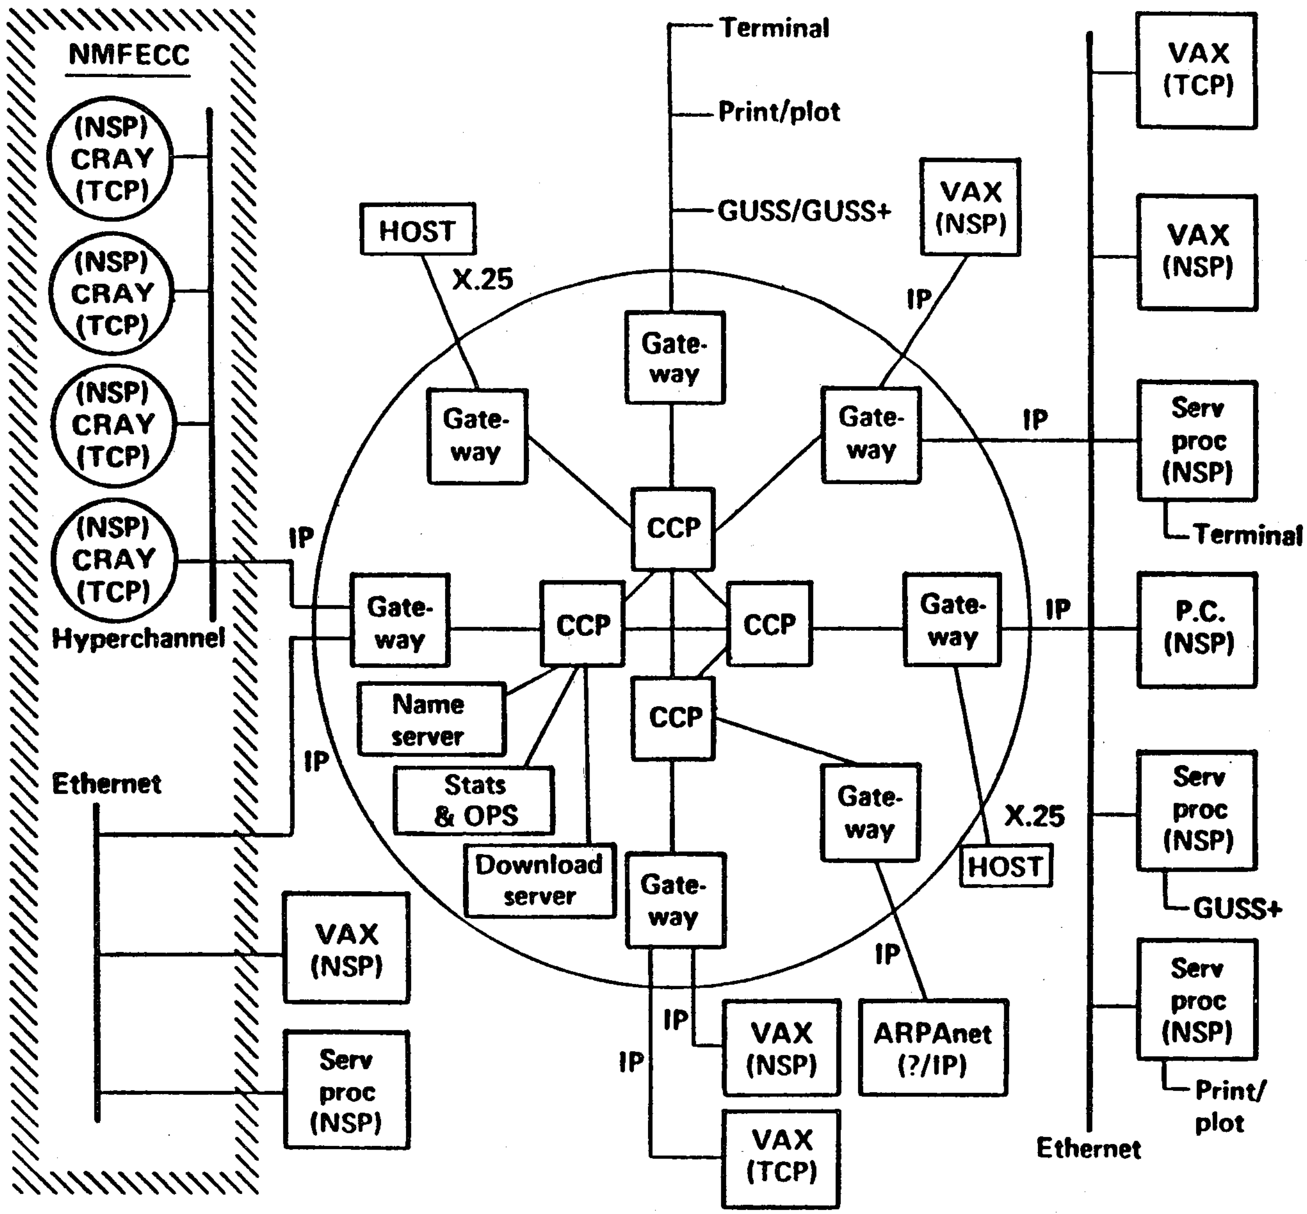

FIGURE III

MFEnet II FIVE YEAR PLANNING RELATED TO THE ISO OSI MODEL

	ISO MODEL	ISSUE	LOCATION	EXAMPLES		5 YEAR PLAN
7	APPLICATION	FUNCTION		TELNET NETIN FTP TELL/MESSAGE NETTY NETOUT FTAM		FUNCTIONALITY
6	PRESENTATION	DATA FORMAT		?		
5	SESSION	SESSION	HOST	SOCKETS		INTER-OPERABILITY
4	TRANSPORT	END-TO-END CONTROL		NSP TCP TP-4		
3.5	INTERNET	ADDRESSING AND ROUTING		IP	X.75	
3	NETWORK			NRP		CONNECTIVITY
2	LINK	POINT-TO-POINT RELIABILITY	NETWORK	DDCMP HDLC SDLC	X.25	
1	PHYSICAL	HARDWARE		V.35 RS-232 RS-422		

FIGURE IV

MFEnet II

Full Backbone Configuration

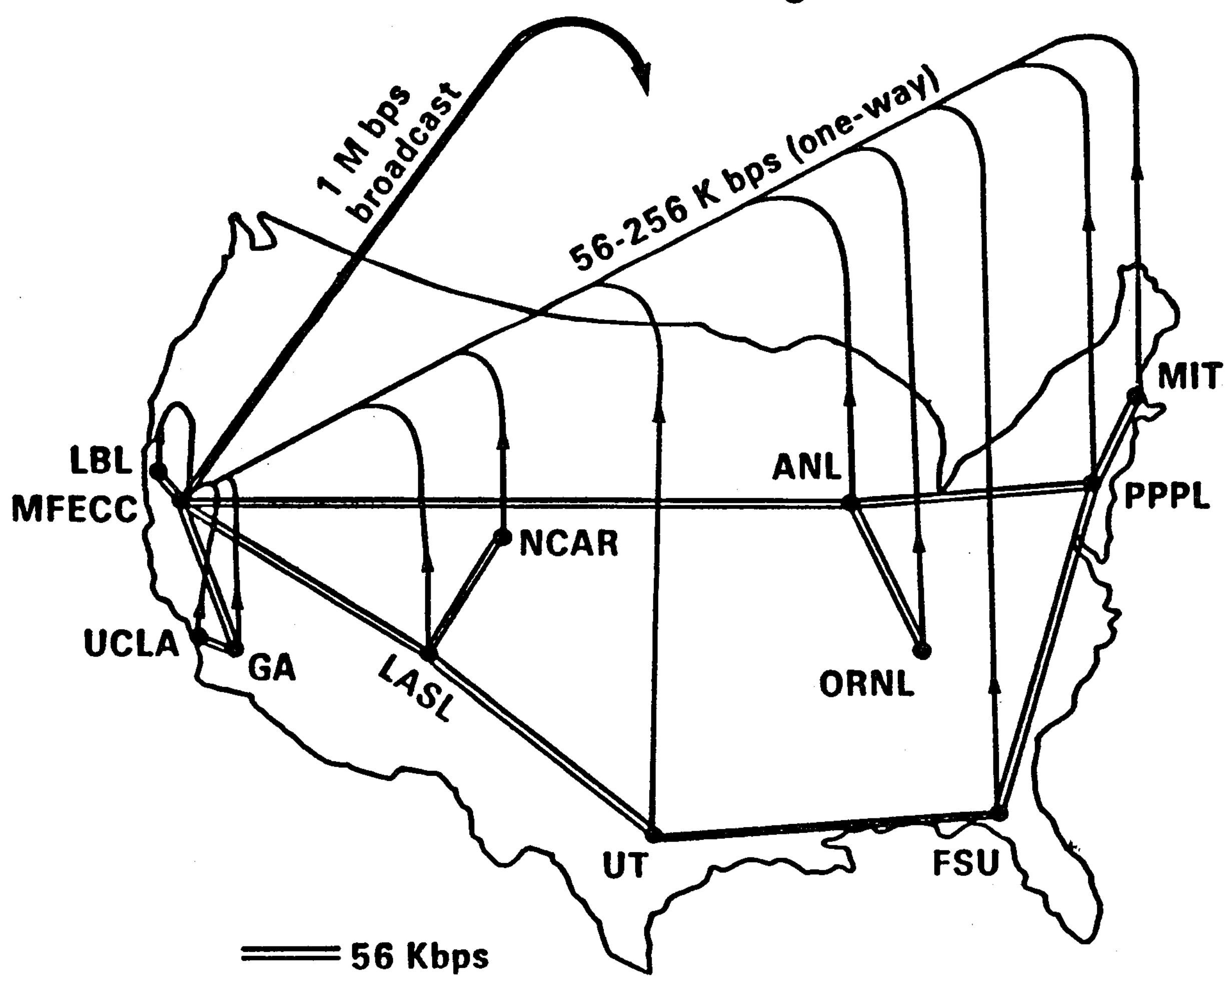

FIGURE V

International Fusion Program

Europe (approx. 65 total sites)
Japan (approx. 45 total sites)

Principle Sites - Ongoing Research

Germany
* Garching - IPP, Max Plank Institute

Julich KFA
Kuhlsruhe - KFK

UK
* Culham/Harwell - JET, Culham Laboratory

France
Cadarache - CEN

Italy
Frascati - Frascati Energy Research Center - ENEA
Padova - IGICGR - ENEA

Austria
Vienna - IAEA Headquarters

Switzerland
Lausanne - CRPP, EPFL
Geneva - CERN

Japan
Kyoto - Kyoto Univ, Uji
Nagoya - IPP
* Naka - JAERI
Tokai - JAERI
KEK

(* Largest Facilities)

In addition, there is in planning a major design effort for ETR. This will involve Europe, Japan, USA, and USSR and extend from 1988 through 1990. Approximately 50 engineers/community for a total of 200 would be involved. The design would be a joint effort and would require extensive exchange of drawings, calculations, CAD/CAM data, etc. The principle sites for design work and coordination would be:

Europe - IPP Garching (NET Design center)
Japan - Naka (JAERI)
USA - distributed
USSR - probably Moscow
Vienna - IAEA Headquarters

DOE/ER-0341

U.S. DEPARTMENT OF ENERGY
OFFICE OF ENERGY RESEARCH

ESNET

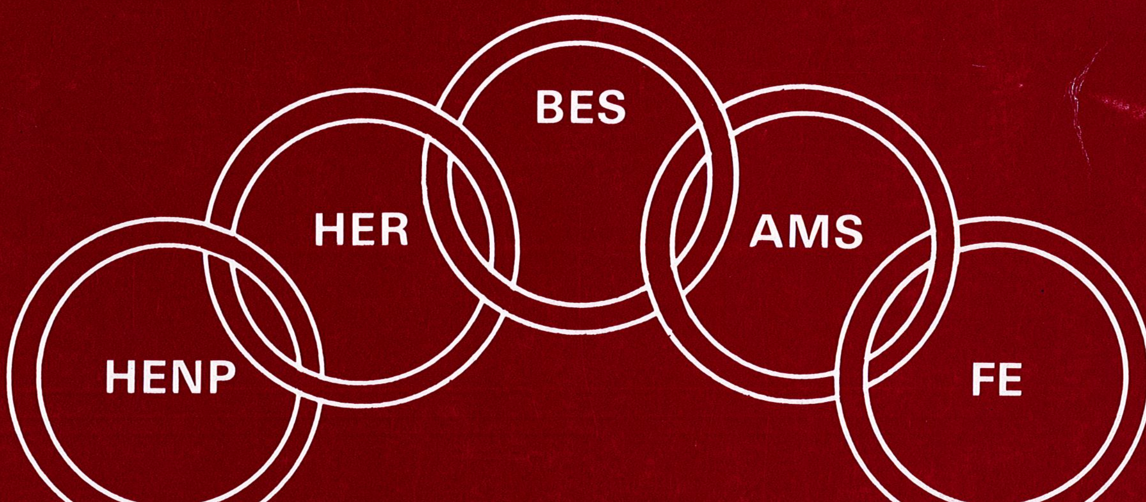

June 1987

DOE/ER-0341

U.S. DEPARTMENT OF ENERGY
OFFICE OF ENERGY RESEARCH

ESNET

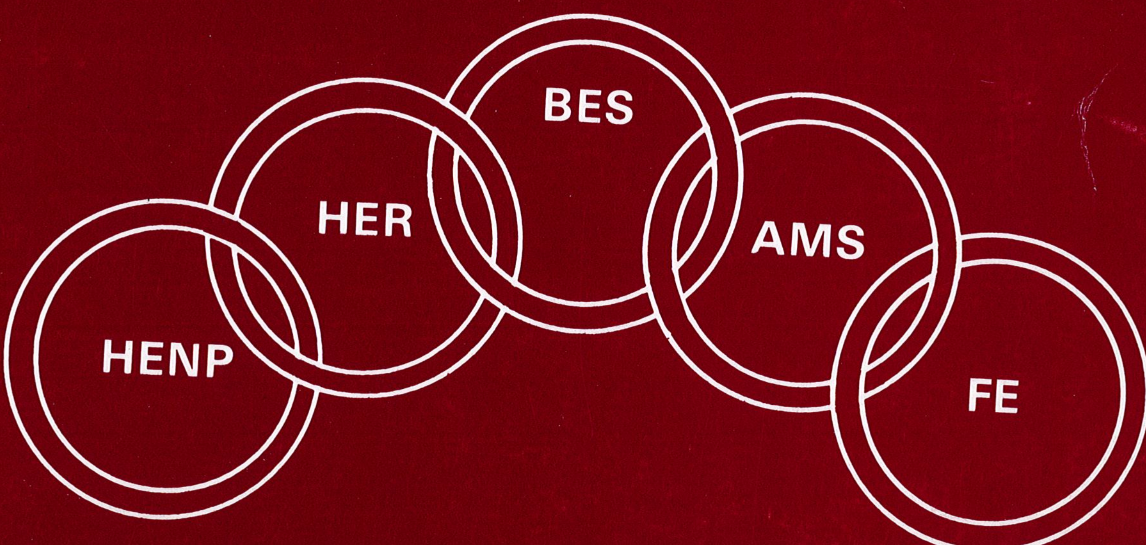

June 1987

UNITED STATES
DEPARTMENT OF ENERGY
WASHINGTON, D.C. 20545

ER-7

APPLIED MATHEMATICAL SCIENCES

NETWORK REQUIREMENTS

AMS PROGRAM NEEDS FOR ESNET

The AMS community surveyed the networking requirements of the Applied Mathematical Sciences (AMS) Research Programs. Sites surveyed include: AMES, ANL, BNL, CIT, FSU, LBL, LANL, LLNL, NYU, ORNL, Pacific Northwest, Sandia, University of Illinois. (See attachments). Below you will find the specific questions that AMS researchers were asked to address accompanied by summaries of their responses.

Identify sites with which you collaborate:

All National Laboratories, most major universities, commercial vendors, specific DOD Laboratories, CERN, and Universities at Edinburgh, Hamburg, Saclay.

Identify networks to which you require access:

ARPA/Milnet, BITNET, CSNET, HEPNET, MFENET, NSFNET inc. associated regional networks, TYMNET, UUCP, USENET (It is generally understood that ESNET will encompass the heretofore HEPNET and MFENET).

Identify needed networking services:

All sites utilize virtual terminal, remote file transfer, mail, remote job entry, bulletin board, and phone/talk services. Projections exist for a need to support distributed computing, including: distributed operating systems, distributed graphics, remote procedure calls, and network file systems. A subpanel of the ESNET Steering Committee, authored by Dunning, Greenwald and Loken, has created a document "Desired Functional Characteristics of ESNET". This document can be endorsed as containing the start up requirements for a network intended to meet today's networking needs of the Applied Math programs. Note its similarity in content to the attached input from AMS researchers. Please also reference the attached LBL Report "The Effect of Distributed Computing Technology on Wide Area Network Capacity Requirements." Note that known technology trends will increase the required bandwidth of ESNET during the upcoming five years by a few orders of magnitude, of which at least one order of magnitude will need to be met by increased bandwidth. The proposed satellite broadcast mechanism provides the needed communication path to satisfy the networking needs resulting from: (1) the future implementation of large distributed databases which must be concurrently updated, (2) the need to transfer very large files without interfering with interactive traffic, and (3) backup communications path for critical interactive traffic.

Define needed bandwidth:

Concerns surrounding bandwidth include not only ubiquitous access and the total capacity of the network in terms of bits per second, but also the network's ability to handle response time critical data concurrently with throughput critical data. Specific bandwidths requested by the major sites include 56 Kilobits and 1.5 Megabits per second.

Identify utilized protocols:

The protocols most frequently named by the AMS community include: TCP/IP, MFENET, and DECNET suites. There was mention of HASP, ASYNC, and movement toward ISO. The selection and implementation of IP as the transport mechanism for ESNET is an important step in facilitating ubiquitous and cost effective connectivity throughout our user community. Such positioning will allow ESNET to migrate to ISO protocols with minimal effort and delay, a critical element in allowing full interoperability among not only DOE resources but those within universities and agencies with whom we collaborate.

In closing I should note that AMS users rely heavily on networking to access Supercomputer resources and to collaborate with each other. Thus the IP implementation of ESNET is expected to enhance our ability to meet programmatic needs.

Attachments provided

Oak Ridge National Laboratory

Requirements

In reponse to your request for network information dated Nov. 13, 1986, here is a summary of our networking needs:

1. sites with which we collaborate
   DOE labs: LLNL (incl. MFE), ANL, LANL, LBL, BNL, SNLL, most major universities, numerous commercial vendors, and military labs and bases.
2. networks to which we require access
   DDN (MILNET and ARPANET), BITNET, USENET, CSNET, NSFNET, MFENET.
3. needed networking services
   mail, remote login, file transfer, task-to-task, network file systems, bulletin boards, talk/phone
4. needed bandwidth
   56Kbs, T1 (desired)
5. protocols used
   TCP/IP, NFS, UDP, telnet, ftp, smtp, DECNET

Most of these services are presently available to us through our DDN connection, but higher bandwidth and better performance are desired. Access to the MFE machines should be enhanced to support TCP/IP access, eliminating the need for TCP to NSP software development.

(3) "networking services"

Networking services required will certainly be mail, file shipping, and remote supercomputing. In addition, there will be a limited need for secure transmission of classified messages and data. Today, analysis of many computational results requires sophistocated graphics. Given this requirement, it will be a challenge for future networks to adequately support remote supercomputing.

(4) "bandwidth requirements"

The present dual 56Kbit link from Los Alamos to MFE is probably a bare minimum. These requirements are closely coupled to the services that you plan to support. If any distributed applications are planned (e.g. debugging, graphics) then the bandwidth requirements might be significantly greater than a predominantly "file shipping" operation.

(5) "protocols used etc."

It is our feeling that ESNET should initially conform to the DDN protocols (TCP/IP) and migrate to ISO protocols. With this capability, the network would have increased interoperability between various other networks and perhaps lower operating costs.

If you have further questions or would like any of the points expanded upon please let me know. Electronic mail address is abw@lanl and phone is 505-667-7028. I hope this is not too late to be of assistance.

FSU - NETWORK REQUIREMENTS

1) The sites with which we here at the Supercomputer Compututations Research Institute of Florida State University collaborate are many and far ranging since this is a multi-disciplinary effort. On a regular basis, we have requirements for data communications to the following institutions

DOMESTIC:

FNAL, BNL, ANL, LBL, LLNL, SLAC, NCAR, SDSC, Los Alamos, ORNL, CEBAF, ORAU, Naval Research Lab, Univ Rochester, MIT, Princeton, Univ Alabama, Lehigh Univ, Univ Florida, Univ Illinois / Chicago, Univ Minn, Northwestern, Univ Connecticut, Univ Colorado, Univ Arizona, Univ Houston, NYU, GWU, and Univ Kentucky.

FOREIGN:

CERN, Univ Edinburgh, Karlsruhe, Imperial College of Science and Technology, Univ Hamburg, Saclay, Hebrew Univ and others by our visitors.

2) We currently have access to MFENET, BITNET and TYMNET and use them on a regular continuing basis. Access to the ARPA-Internet, the NSFnet backbone with attendant regional and consortium networks is desired and should be operational by the end of February. We are planning connection to the MFENET II (ESNET?) sometime next summer. Connection to HEPnet is also desired, although this may be the same as MFENET II in the long run.

3) The services which we feel are needed in the upgraded network(s) of the future for access to our resources (Cyber 205, ETA-10) by others and to other resources by us include:

Remote login to another host with sufficiently transparent access across the connection to allow for graphics work and full-screen editor work from a terminal that normally supports such activity.

Remote file retrieval / delivery from one host to another, with options for handling both character information and binary data.

FSU - NETWORK REQUIREMENTS

Remote job submission / output retrieval where the user is dealing with the input and output queues of the remote host rather than a file store as was the case in the last item. Options required include complete transparency on the output for graphics data.

Electronic mail to the network / from the network / redirected to another mailbox (potentially, on another host across the network) / with a rich feature set including message confirmations. A consistency of address formulations is also desirable, but probably not attainable until the ISO OSI standards are implemented, if then.

Provision of remote procedure calls in the network would allow the community to add anything else that was truly needed by a small number of people.

4) The bandwidth required changes over time (almost always up) as people learn more and more of what is possible. We currently have 9600 baud service on MFEnet, BITNET and TYMNET. There are definitely requirements for more but how much more is not certain. The link to CERN, currently being implemented, was sized at 56 / 64 Kbaud to provide good response time for shipping files back and forth. The SURAnet link, currently being implemented, for the NSFnet connection will be at 56 Kbaud and also provide for access to the Internet community. When the ETA-10 comes on line as a production system in the Fall of 1987, the demand for access will surely swamp the existing networks, if nothing is done.

5) We are currently using / supporting the following protocols:

MFENET I, from NMFECC
BITNET, JES2/NJE from an IBM MVS system
TYMNET, X.25
DECNET, to other VAXes, both on and off campus
TCP/IP, for the SURAnet/NSFnet, the ARPA-Internet and planned for MFENETII
HASP, for remote job entry station access
ASYNC, dial-up for people not connected to one of the networks.

Los Alamos National Laboratory

Requirements

I have talked with several people about issues relating to the Energy Sciences Network (ESNET). Their comments on your five points are summarized below.

(1) "sites with which you collaborate"

It would be very difficult to compile a complete list of all the sites with which Los Alamos scientists collaborate. Such a list would certainly include all DOE Laboratories and Offices, all NSF Supercomputer Centers, all DOE production plants, all University of California schools and all New Mexico Universities.

In addition to these I would include University of Texas - Austin, Cal Tech, Stanford, Cornell, Carnegie-Mellon, Brown University, RPI, Simon Fraser, University of British Columbia, University of Arizona, Texas Tech, University of Colorado - Denver and Boulder, University of Wisconsin - Madison, NYU, and MIT.

(2) "networks to which you require access"

It is difficult to guess what networks we will require access to in the future, but adherence to a protocol standard would make future access simpler. The DOE is involved with several networking projects (e.g. NWCNET and OPMODEL) and there are alot of existing networks (see October 1986 issue of Communications of the ACM or 28 February 1986 issue of Science magazine). It would seem reasonable to fold as much of ESNET into existing efforts as possible.

There has been a considerable increase in the need for open computing cycles here, thus we expect that Los Alamos scientists will require access to all DOE and NSF supercomputer centers (in addition to local resources). The networks that we have direct or indirect access to at the moment, and will require continued access to, are DDN (Arpanet/Milnet), BITNET, CSNET, TELENET, NMFENET, NWCNET (not yet fully operational), TECHNET (a New Mexico, statewide communications network), NSFnet (under development), UUCP, and USENET.

Argonne National Laboratory

Requirements

this is in response to your request for information on our networking plans. We currently make extensive use of ARPAnet/MILnet, MFEnet, BITnet, CSnet and TYMNET. Nobody in our division uses HEPnet. In addition, we are attempting to set up a high-performance link to the University of Illinois at Urbana. This will be used in collaborative research at traffic levels that cannot be met by existing links. In general, we are interested in any new network serving sites we cannot now reach. We support the idea of centralized funding for gateways to such networks. I suspect that none of the other connections is exotic enough for you to pursue funding at a higher level.

As discussed in the appended extract from a draft of a white paper (prepared by Jim Boyle of our staff) that may be published in the Congressional Record, we believe that in the future we will be interested in distributed computation over wide area networks. This would require greater network performance than now exists, which is an issue you might wish to consider.

NETWORK FUNCTIONAL REQUIREMENTS

We believe that from the users' point of view that there are four major types of functional requirements for computer networks. They are:

- Electronic mail
- File transfer
- Terminal access
- Distributed computation

We believe that users require these services not just within their local network, but within hierarchies of local networks and within wide area networks, including world-wide ones. In particular, access to these functions must be transparent through network gateways.

If satisfactory service is to be provided, each of these functions induces its own requirements on the design of networks and their interconnection mechanisms. In the following discussion, we classify these requirements in respect to the network properties:

- Latency: the delay in passing a message across a network
- Speed: the rate at which characters can be transmitted
- Bandwidth: the total capacity of the network per second

Electronic Mail

Electronic mail involves the sending of a message to a user on a machine, possibly one that is several networks away from the sender's machine. From the user's point of view, the message should arrive in a timely fashion, say within a few minutes if all required networks are operating. When a required network is down, the message should be queued for delivery when it is again available. Information about the disposition of the message (for example, return receipt requested) may be desired. Unlike a telephone call, receipt of the message must take place without any action on the part of the recipient. Electronic mail messages are typically small. However, they provide a convenient mechanism (because simultaneous interaction of the sender and recipient is not required) for the exchange of manuscripts or computer programs, whose size may be 100,000 bytes or more. The user requires convenient access to a "Yellow Pages" server that lists electronic mail addresses.

Existing electronic mail schemes, such as those used on the ARPAnet, are basically satisfactory. Difficulty of interchange and addressing across dissimilar networks are currently weak points, but emerging mail standards can be expected to solve many of these problems independently of network advances.

Electronic mail demands only limited speed and responsiveness from the network. However, to be useful it does require a store-and-forward capability somewhere in the network. For electronic mail to be useful, the networks used must be reliable in the sense of being available without long outages.

File Transfer

File transfer involves a user initiating a request to send a file or files to, or to receive a file or files from, a remote computer. From the user's point of view, he would like the files to be transferred in real time and as quickly as possible. He would like easy access to the remote system, while at the same time he would like to be assured that his files are protected from unauthorized access by others. To be effective, file transfer requires some limited terminal-like access to the remote system, in order to find the directory and file to or from which transfer is to take place. The volume of data being transferred can be quite high, 10 megabytes or more, and such volumes need to be transferred in at most a few minutes. Of course, the data must arrive without error, and with suitable character translations having been made. At the same time, binary or other machine-specific data must not be corrupted in the transfer.

Existing file transfer mechanisms (for example, ftp on the ARPAnet, and especially rcp among Unix systems) are basically satisfactory. Higher speed and greater network bandwidth would improve usability.

File transfer demands high network speed to minimize the time a user must spend effecting the transfer. Network bandwidth must be adequate to handle the anticipated number of simultaneous file transfers. Latency is not a serious problem except when interactive listing of directories or files is taking place.

Terminal Access

Terminal access involves using a remote computer from a terminal or workstation as if it were a local resource. From the user's point of view, he should not be able to tell that the computer he is using is thousands of miles away. Currently, a typical "message" from a user to the remote resource in terminal access mode involves a few characters--in many cases, a single character. Messages from the remote resource to the terminal are larger on the average, although single character messages will occur when echoing in full-duplex mode. In the future, larger messages may be more common, for example, for graphical output from the remote resource.

Existing terminal access services, such as telnet on the ARPAnet and especially rlogin among Unix systems, are basically satisfactory within local area networks. Currently, however, wide area networks do not adequately meet the requirements discussed in the next paragraph, resulting in user dissatisfaction.

Terminal access currently requires only relatively low network speed (tens of Kbaud) and limited bandwidth. However, it sets stringent limits on the latency of the network. Many applications must run full-duplex (or generalizations thereof, such as graphical response to user input from a mouse). Network latency of even several tenths of a second can seriously impair usability. Achieving low latency may preclude the use of satellite links in wide area networks, and places stringent demands on latency in gateways. Similarly, sufficient bandwidth must exist in the network if low latency is to be maintained while other high-bandwidth functions are taking place. When connecting LANs by WANs, it may be necessary to think in terms of two parallel networks: a low latency one for terminal access, which would use overland and underwater links, and a higher-latency one for the other functions using satellite links.

Distributed Computation

Distributed computation involves running a multiprocessing user program on two or more machines on a network simultaneously. We expect such applications to be of increasing importance over the next few years.

A few initial applications are already in place. One is remote access to file systems (for example, via the Network File system, NFS). In this application, a file system on a remote computer can be mounted and made to appear to be a local file system on the user's computer. Remote files are accessed by sending individual blocks of the file over the network in real time when required, rather than by transferring the entire file at one time. We believe such remote file system use will become increasingly common, with files shared not only nationally but internationally.

Another distributed application that is beginning to be put in place is remote display of graphical output, in which graphical output is sent from a large supercomputer via the network to a workstation for display. We believe that such applications will grow rapidly. They may be driven in part by "electronic publishing", in which magazines and newsletters, complete with photographs and graphics, would be distributed via network and displayed on a reader's workstation. We think use of distributed text editors, and text composition systems, for example, will also soon be common.

Finally, within the next one to two years distributed simulations and other application-oriented computations will become common. Such distributed computations will require the relatively frequent exchange of messages over the network to maintain synchronization among the distributed processes.

Existing distributed computations are either prototypes or special-purpose (for example, NFS). Thus, the adequacy of present networks to support distributed computation is largely unknown.

Distributed computation can be expected to place severe demands on all three network parameters: latency, speed, and bandwidth. Remote file access requires substantial bandwidth and low latency, since the file may be used interactively. Graphical output makes similar demands, especially in regard to latency, because the user may need to interact with the program running on the supercomputer that produces the output. Distributed simulations and other application programs will require high speed and low latency, and possibly high bandwidth, to effect synchronization among parallel processes via message passing.

Complicating these requirements of individual distributed applications is the fact that many individuals and groups will be trying to perform them simultaneously, raising the overall network bandwidth requirements even higher. It seems clear that use of NFS over national distances would require apparent Ethernet bandwidth for each such use, generating a requirement for a WAN of several times Ethernet speed.

Finally, we remark that "combining networks" may play a role in satisfying these demands, although much research would be required. Combining networks would attempt to avoid multiple transmission of the same packets, something that must occur millions of times a day on the ARPAnet, when messages are distributed to mailing lists.

HEALTH & ENVIRONMENTAL RESEARCH

NETWORK REQUIREMENTS

The Office of Health and Environmental Research (OHER) supports a considerable number of projects that require supercomputer usage and/or interlaboratory computer networking. I have asked three scientists to represent OHER in various aspects of this activity. Chuck Watson, an expert in computerized database design, has been detailed here from PNL to develop an integrated approach to storage and analysis of OHER supported toxicological studies. I have asked him to coordinate his efforts with you. Walter Goad, director of GENBANK at LANL, is an active participant on the ESNET steering committee representing the OHER research community. George Duda, an OHER staff member, is responsible for coordination of MFE computing time for OHER. They will be developing an overview of the OHER long range plan for computer usage which will aid you in your planning activities.

Until the comprehensive plan is in place, the following items related specifically to ESNET requirements should be useful. These items represent a summary of the attached earlier correspondence from Watson, Goad and others to your Office.

Sites With Which OHER Researchers Collaborate

OHER researchers collaborate with fellow researchers in all National Laboratories, and most major universities. It is important to note that one major laboratory, the Inhalation Toxicology Research Institute in Albuquerque, New Mexico is not connected to any computer network. Foreign collaboration includes frequent contact with researchers at the European Molecular Biology Laboratory in Heidelberg, Federal Republic of Germany, and the DNA Data Bank of Japan (DDBJ) at the National Institute of Genetics, Misima, SizuokaKen, Japan, which is about half way between Tokyo and Nayoga on the island of Honshu.

Network Which OHER Researchers Access

The OHER research community currently makes use of BITNET, ARPA/Milnet, MFENET, the NSFNET and its associated regional networks, and TYMNET.

Needed Network Services for OHER Researchers

At a minimum, all sites need virtual terminal access, remote job entry and file transfer, mail, and bulletin board services. The largest growth areas for the OHER research community will be in distributed graphics and distributed database technologies. These will be especially important in the fields of molecular dynamics and protein sequence analysis.

Anticipated Bandwidth Required by OHER Researchers

At a minimum, all OHER researchers need 9600 baud access. Distributed graphics for selected researchers will require orders of magnitude more bandwidth.

Network Protocols Utilized by OHER Researchers

The DECNET has the most users in the OHER research community.

I want to confirm my support for the consolidation of the various DOE research networks into ESNET. The OHER research community can only benefit from this orderly upgrade of services. The reduction in parallel communications lines and software development efforts is very welcome. I appreciate this opportunity to provide input to your planning process. Please direct questions to Chuck Watson.

HYDROLOGY REQUIREMENTS

Requirement	Current	5 Years	Locations
Networking-Mail	Yes*	Yes*	All National Lab Sites, (especially PNL, ORNL, LANL, LLNL, SRL)
Scientific Collaboration (exchange of information among cooperating groups)	Yes* This as a major element in ERD's hydrology programs	Yes*	PNL to other lab sites SRL, ORNL, LLNL, LANL
File Transfers	No	No	--
Video Conferencing	Yes*	Yes* (increased)	All lab sites
Supercomputer Access	Limited	Yes (if they do not have own system)	PNL, ORNL to LLNL
Library Access	Yes	Yes	Technical Information Center, ORNL
Graphics	No	none to very limited	--

*NOTE: There are many possible links to university collaborators, but these are very diffuse and change in a few years.

OHER MFE Super Computer Usage

Table 1. FY 87 allocation of MFE Central Resource Units (CRU) tabulated by OHER division

division	CRU	percent	division name
HERD	535	37	Health Effects Research
PTRD	516	35	Physical and Technological Research
HHAD	260	18	Human Health and Assessments
ERD	80	5	Ecological Research
reserve	69	5	
total	1460		

Table 2. FY 87 OHER allocation of MFE CRUs tabulated by geographic location

State	laboratory	PTRD	ERD	HERD	HHAD	TOTAL
Calif	U C Berkeley			80		80
Calif	LBL	30			260	290
Calif	USC	20				20
Calif	Stanford		6			6
Calif	LLNL	30				30
Fla	Miami		5			5
Fla	U. FLA	36				36
Ga	Skidaway		30			30
Illinois	ANL	40		40		80
Miss	U Missouri		15			15
NM	LANL	25				25
NY	U Rochester			10		10
NY	Columbia	100				100
NY	BNL	9		170		179
NY	NYU			150		150
Pa	Penn State	90				90
Tenn.	ORNL	5	18	80		103
Wash	U Wash.		6			6
Wash	PNL	131		5		136
Reserve						69
	totals	516	80	535	260	1460

institution with ESNET. I will have to expand the table to reflect this after the contracts are awarded and I am able to talk with the PI's about their network access.

The motivational needs can be addressed via coercion. DOE can (and will) require that management reporting be done via the computer. DOE can also insist that research findings be sent to PNL via computer. However, effective use of a computer network can not be legislated. People need to be trained, they need to be shown that sending information to a remote colleague is not difficult. Most importantly, people need to develop a trust that their data will be acknowledged and used by responsible scientists. I plan to participate in the radon program kick off meeting in March to start the motivational aspects of my task. Later, I will be visiting key sites within the program.

We already have some information about the data flow within this radon network. LBL is a major player in the existing radon project. Since LBL is on ESNET, it is natural to center our network efforts at LBL. ITRI is a major player in the new thrust. Highest priority should be given to acquisition of a dedicated phone line connecting ITRI to the network, probably via LASL.

Two other sites will generate considerable data, EML in New York, and PNL in Richland, both are current members of ESNET

The table summarizes the situation for FY 87. However, we know that additional projects will be funded in FY 88 and beyond. At this point, we do not know the geographic location of these new projects. It is safe to assume, nevertheless, that at a minimum, dial up access to ESNET will be convenient from any new radon program laboratgory.

Table 1 - Radon Research Laboratories

Laboratory	Title
DOE Headquarters	Program Management
PNL	Mechanism of Injury to Critical Cells of the Respiratory Tract from Inhaled Radon Daughters
	Critical Bibliography on the Health Effects of Radon
ITRI	Program dsecription document
	Radiation Dose and Injury to Critical Cells of the Respiratory Tract from Inhaled Radon Progeny
LBL	Indoor Air Quality Database
	Behavior of Radon 222 and its Decay Products in the Indoor Environment
	Studies of the Molecular Damage in Neoplastic Transformation caused by Alpha Particle Radiation from Radon Daughters
ANL	Title Unknown
ORNL	
EML	Instrumentation Development
BPA	Residential Weatherization Program - Radon Monitoring
U of New Mexico	Lung Cancer Epidemiology
Inst. of Mines, Socorro, NM	Title Unknown
Texas A&M	Aerosol Microphysics of Indoor Radon
NY U Med Center	Lung Cancer Risk from Inhalation of Radon and Other Pollutants in Rats
	Measurement and Apportionment of Radon Source Terms for Modeling Indoor Environments
U of Illinois	Study of the Atmospheric Chemist of ^{218}Po
Dept of Interior Denver, Co	Gas Transport in Soils and Its Relation to Radon Availability
	Investigations of the Relation Between Aeroadiometric Bi-214 Signatures and Soil Radioactivity, Radium, and

	Radon Content	
Yale	Determination of ^{222}Rn flux in soils based on ^{210}Pb and ^{226}Ra dsequilibrium	
Penn State U	Generation and Mobility of Radon in Soil	
Colorado State U	Effects of Vegitation on Radon Transportation Processes in Soil	

Table 2
Radon Program
Computer Network

Geographic Distribution
Fiscal Year 1987

Currently funded:	NETWORK ACCESS
DOE Headquarters Germantown, MD	Dial up to Univ of Maryland
PNL Richland Washington	Current ESNET member
EML New York City	Current ESNET member
BPA Portland, Oregon	Can get to ESNET via PNL
U of New Mexico Albuquerque, NM	dial up to LASL or ITRI if leased line installed between ITRI and LASL.
Ins. Mining & Tech. Socorro, NM (~50 miles South of Albuquerque)	dial up to LASL or ITRI if leased line installed between ITRI and LASL.
ANL Chicago, Il	Current ESNET member
New activities FY 87:	
Texas A&M Univ. College Station, TX	dial up to U of Texas ?
New York U Med Center New York City	dial up to EML ?
Univ. of Illinois Urbana Il	Current ESNET member
Dept of Interior Denver Co. Reston VA ? (Need to talk to PI re involvement from Reston.)	New ESNET thrust should add node near Denver in CY 87 Reston could dial up to U of Maryland.
ITRI Albuquerque, NM	currently dial up to LASL I recommend that a leased line be installed between ITRI and LASL.
Yale University New Haven, Conn.	dial up to ?
Penn State U University Park, PA	Current ESNET member
Colorado State Univ Ft. Collins, CO	New ESNET thrust should add node near Denver in CY 87
ORNL Oak Ridge, Tn	current ESNET member

HEALTH & ENVIRONMENTAL RESEARCH

NETWORK REQUIREMENTS

The Office of Health and Environmental Research (OHER) supports a considerable number of projects that require supercomputer usage and/or interlaboratory computer networking. I have asked three scientists to represent OHER in various aspects of this activity. Chuck Watson, an expert in computerized database design, has been detailed here from PNL to develop an integrated approach to storage and analysis of OHER supported toxicological studies. I have asked him to coordinate his efforts with you. Walter Goad, director of GENBANK at LANL, is an active participant on the ESNET steering committee representing the OHER research community. George Duda, an OHER staff member, is responsible for coordination of MFE computing time for OHER. They will be developing an overview of the OHER long range plan for computer usage which will aid you in your planning activities.

Until the comprehensive plan is in place, the following items related specifically to ESNET requirements should be useful. These items represent a summary of the attached earlier correspondence from Watson, Goad and others to your Office.

Sites With Which OHER Researchers Collaborate

OHER researchers collaborate with fellow researchers in all National Laboratories, and most major universities. It is important to note that one major laboratory, the Inhalation Toxicology Research Institute in Albuquerque, New Mexico is not connected to any computer network. Foreign collaboration includes frequent contact with researchers at the European Molecular Biology Laboratory in Heidelberg, Federal Republic of Germany, and the DNA Data Bank of Japan (DDBJ) at the National Institute of Genetics, Misima, SizuokaKen, Japan, which is about half way between Tokyo and Nayoga on the island of Honshu.

Network Which OHER Researchers Access

The OHER research community currently makes use of BITNET, ARPA/Milnet, MFENET, the NSFNET and its associated regional networks, and TYMNET.

Needed Network Services for OHER Researchers

At a minimum, all sites need virtual terminal access, remote job entry and file transfer, mail, and bulletin board services. The largest growth areas for the OHER research community will be in distributed graphics and distributed database technologies. These will be especially important in the fields of molecular dynamics and protein sequence analysis.

Anticipated Bandwidth Required by OHER Researchers

At a minimum, all OHER researchers need 9600 baud access. Distributed graphics for selected researchers will require orders of magnitude more bandwidth.

Network Protocols Utilized by OHER Researchers

The DECNET has the most users in the OHER research community.

I want to confirm my support for the consolidation of the various DOE research networks into ESNET. The OHER research community can only benefit from this orderly upgrade of services. The reduction in parallel communications lines and software development efforts is very welcome. I appreciate this opportunity to provide input to your planning process. Please direct questions to Chuck Watson.

HYDROLOGY REQUIREMENTS

Requirement	Current	5 Years	Locations
Networking-Mail	Yes*	Yes*	All National Lab Sites, (especially PNL, ORNL, LANL, LLNL, SRL)
Scientific Collaboration (exchange of information among cooperating groups)	Yes* This as a major element in ERD's hydrology programs	Yes*	PNL to other lab sites SRL, ORNL, LLNL, LANL
File Transfers	No	No	--
Video Conferencing	Yes*	Yes* (increased)	All lab sites
Supercomputer Access	Limited	Yes (if they do not have own system)	PNL, ORNL to LLNL
Library Access	Yes	Yes	Technical Information Center, ORNL
Graphics	No	none to very limited	--

*NOTE: There are many possible links to university collaborators, but these are very diffuse and change in a few years.

OHER MFE Super Computer Usage

Table 1. FY 87 allocation of MFE Central Resource Units (CRU) tabulated by OHER division

division	CRU	percent	division name
HERD	535	37	Health Effects Research
PTRD	516	35	Physical and Technological Research
HHAD	260	18	Human Health and Assessments
ERD	80	5	Ecological Research
reserve	69	5	
total	1460		

Table 2. FY 87 OHER allocation of MFE CRUs tabulated by geographic location

State	laboratory	PTRD	ERD	HERD	HHAD	TOTAL
Calif	U C Berkeley			80		80
Calif	LBL	30			260	290
Calif	USC	20				20
Calif	Stanford		6			6
Calif	LLNL	30				30
Fla	Miami		5			5
Fla	U. FLA	36				36
Ga	Skidaway		30			30
Illinois	ANL	40		40		80
Miss	U Missouri		15			15
NM	LANL	25				25
NY	U Rochester			10		10
NY	Columbia	100				100
NY	BNL	9		170		179
NY	NYU			150		150
Pa	Penn State	90				90
Tenn.	ORNL	5	18	80		103
Wash	U Wash.		6			6
Wash	PNL	131		5		136
Reserve						69
	totals	516	80	535	260	1460

institution with ESNET. I will have to expand the table to reflect this after the contracts are awarded and I am able to talk with the PI's about their network access.

The motivational needs can be addressed via coercion. DOE can (and will) require that management reporting be done via the computer. DOE can also insist that research findings be sent to PNL via computer. However, effective use of a computer network can not be legislated. People need to be trained, they need to be shown that sending information to a remote colleague is not difficult. Most importantly, people need to develop a trust that their data will be acknowledged and used by responsible scientists. I plan to participate in the radon program kick off meeting in March to start the motivational aspects of my task. Later, I will be visiting key sites within the program.

We already have some information about the data flow within this radon network. LBL is a major player in the existing radon project. Since LBL is on ESNET, it is natural to center our network efforts at LBL. ITRI is a major player in the new thrust. Highest priority should be given to acquisition of a dedicated phone line connecting ITRI to the network, probably via LASL.

Two other sites will generate considerable data, EML in New York, and PNL in Richland, both are current members of ESNET

The table summarizes the situation for FY 87. However, we know that additional projects will be funded in FY 88 and beyond. At this point, we do not know the geographic location of these new projects. It is safe to assume, nevertheless, that at a minimum, dial up access to ESNET will be convenient from any new radon program laboratgory.

Table 1 - Radon Research Laboratories

Laboratory	Title
DOE Headquarters	Program Management
PNL	Mechanism of Injury to Critical Cells of the Respiratory Tract from Inhaled Radon Daughters
	Critical Bibliography on the Health Effects of Radon
ITRI	Program dsecription document
	Radiation Dose and Injury to Critical Cells of the Respiratory Tract from Inhaled Radon Progeny
LBL	Indoor Air Quality Database
	Behavior of Radon 222 and its Decay Products in the Indoor Environment
	Studies of the Molecular Damage in Neoplastic Transformation caused by Alpha Particle Radiation from Radon Daughters
ANL	Title Unknown
ORNL	
EML	Instrumentation Development
BPA	Residential Weatherization Program - Radon Monitoring
U of New Mexico	Lung Cancer Epidemiology
Inst. of Mines, Socorro, NM	Title Unknown
Texas A&M	Aerosol Microphysics of Indoor Radon
NY U Med Center	Lung Cancer Risk from Inhalation of Radon and Other Pollutants in Rats
	Measurement and Apportionment of Radon Source Terms for Modeling Indoor Environments
U of Illinois	Study of the Atmospheric Chemist of ^{218}Po
Dept of Interior Denver, Co	Gas Transport in Soils and Its Relation to Radon Availability
	Investigations of the Relation Between Aeroadiometric Bi-214 Signatures and Soil Radioactivity, Radium, and

	Radon Content	
Yale	Determination of ^{222}Rn flux in soils based on ^{210}Pb and ^{226}Ra dsequilibrium	
Penn State U	Generation and Mobility of Radon in Soil	
Colorado State U	Effects of Vegitation on Radon Transportation Processes in Soil	

Table 2
Radon Program
Computer Network

Geographic Distribution
Fiscal Year 1987

Currently funded:	NETWORK ACCESS
DOE Headquarters Germantown, MD	Dial up to Univ of Maryland
PNL Richland Washington	Current ESNET member
EML New York City	Current ESNET member
BPA Portland, Oregon	Can get to ESNET via PNL
U of New Mexico Albuquerque, NM	dial up to LASL or ITRI if leased line installed between ITRI and LASL.
Ins. Mining & Tech. Socorro, NM (~50 miles South of Albuquerque)	dial up to LASL or ITRI if leased line installed between ITRI and LASL.
ANL Chicago, Il	Current ESNET member
New activities FY 87:	
Texas A&M Univ. College Station, TX	dial up to U of Texas ?
New York U Med Center New York City	dial up to EML ?
Univ. of Illinois Urbana Il	Current ESNET member
Dept of Interior Denver Co.	New ESNET thrust should add node near Denver in CY 87
Reston VA ? (Need to talk to PI re involvement from Reston.)	Reston could dial up to U of Maryland.
ITRI Albuquerque, NM	currently dial up to LASL I recommend that a leased line be installed between ITRI and LASL.
Yale University New Haven, Conn.	dial up to ?
Penn State U University Park, PA	Current ESNET member
Colorado State Univ Ft. Collins, CO	New ESNET thrust should add node near Denver in CY 87
ORNL Oak Ridge, Tn	current ESNET member

BASIC ENERGY SCIENCES

NETWORK REQUIREMENTS

Materials Research Needs for Computer Networking

Materials research which encompasses solid state physics, metallurgy and ceramics is concerned primarily with the properties of solids. Traditionally characterized as small science with large numbers of individual investigators an increasing role is being played by larger collaborations and the use of remote facilities. Examples are neutron sources, synchrotron light sources, electron microscopes, and supercomputers.

The Division of Materials Sciences, in the office of Basic Energy Sciences, funds work ranging from individual investigators at universities to intermediate sized groups to large scale facilities at national laboratories. A summary of current work and funding levels is contained in Materials Sciences Programs FY 1986 (DOE/ER-0295). Approximately two-thirds of the budget is spent through DOE laboratories and the remainder at university-based programs, including both laboratories and individual grants. A list of user facilities has been attached to this document.

The main uses of networking at present are by theorists who need access to supercomputers. Recent increases in computing power have been crucial to the development of new techniques for predicting the properties of materials from first principles electronic structure calculations and from simulations on model systems. This has had benefits across the board from the most basic to the most applied research. These supercomputers are located in centers which are often remote from the user. However, with an effective network the user is not aware that the calculations are being performed remotely. In particular, most users are quite content with MFENET in this regard. Therefore, an essential requirement of ESNET is to duplicate all the functions of the current MFENET as far as remote acces to the CRAYs at NMEECC. To put it differently, the network should have a short enough response time so that for editing, job submission etc., the user is not aware that the machine is remote. Users would tolerate some delays for large file transfers. Users would endorse the draft document "Desired functional characteristics of ESNET" by T. D. Dunning, Jr., M. Greenwald, and S. Loken.

In view of the importance of facilities to work in Materials Sciences it can be expected that networking will play a larger role in the future. A typical example involves the collection of a large amount of scattering data from an X-ray or neutron source. This is usually scattered intensity versus scattering angle. Often these data are collected by small computers (PDP-11 or Microvax) and the results stored on floppy disks. Indeed, the scattering geometry is controlled by computer. In addition, the experimenters frequently come from universities or industrial laboratories where they have performed preliminary experiments with laboratory sources. They often return to their home institutions carrying their floppy disks and proceed to analyze their data.

Clearly there is a great potential here for computer networking. First it would mean easier, more accurate data transmission. Ultimately, though, the files may become too large for transport and analysis of remote data bases will become important. However, even the data collection could be controlled remotely and experiments adjusted on line to take advantage of unforeseen developments. This also indicates the need for remote access of software so that preliminary data can be analyzed quickly and experiments altered in progress.

A second example of future uses of networking is the materials preparation center at Iowa State University, Ames, Iowa. Part of this facility is the Materials Referral System and Hotline which accumulates information on the preparation and characterization of materials and makes it available to the scientific community. Clearly, access to such a data base via a network would be a great benefit to workers in materials research.

In summary, the major current use of networking within the Division of Materials Science is remote access of supercomputers. This use is expected to grow as more groups obtain time on the machines and as more sophisticated graphics and editing become available. Increased future use by experimetalists is expected in view of the importance of facilities at remote sites. In addition, use of electronic mail and access to remote data bases can greatly enhance productivity in all materials research programs.

/nw
Attachment

Current Division of Materials Sciences
CRAY USERS

A. Major User Facilities

1. National Synchrotron Light Source	Brookhaven
2. High Flux Beam Reactor	Brookhaven
3. High Flux Isotope Reactor	Oak Ridge
4. Intense Pulsed Neutron Source	Argonne
5. Los Alamos Neutron Scattering Center	Los Alamos
6. Stanford Synchrotron Radiation Laboratory	Stanford
Proposed:	
7. Advanced Light Source	Lawrence Berkeley
8. 6-7 GeV Synchrotron Radiation Source	Argonne

B. Other User Facilities

1. National Center for Small Angle Scattering Research	Oak Ridge
2. Electron Microscopy Center for Materials Research	Argonne
3. Shared Research Equipment Program	Oak Ridge
4. Center for Microanalysis of Materials	Univ. of Illinois
5. Surface Modification and Characterization Collaborative Research Center	Oak Ridge

Communication with Energy Research Supercomputer Centers
Division of Chemical Sciences

Institution	NMFECC Dial up Carrier	NMFECC Direct Access Node	NMFECC Communication Rate
Argonne National Laboratory		Argonne National Laboratory	56
Argonne National Laboratory		Argonne National Laboratory	56
Argonne National Laboratory		Argonne National Laboratory	56
Atlanta University	TYMNet		1.2
Emory University	thru ANL		2.4
Harvard University	TYMNet		1.2
Kansas State University	PARADYNE Modem		19.2
Lawrence Berkeley Laboratory		Lawrence Berkeley Laboratory	56
Lawrence Berkeley Laboratory		Lawrence Berkeley Laboratory	9.6
Louisiana State University	TELENet		1.2
Louisiana State University	TYMNet		1.2
Oak Ridge National Laboratory		Oak Ridge National	48
Radiation Laboratory, University of Notre	TYMNet		1.2
Rice University	TYMNet		1.2
State University of New York/Stony Brook	TYMNet		2.4
University of California	Graduate Student		
University of Colorado	TYMNet (back-up only)	University of Colorado	9.6
University of Florida			1.2
Vanderbilt University	TYMNet		1.2
Yale University	TELENet		

As of 24 March 1987 (incomplete); no significant usage of the computer at Florida State University.

Reference 8

EUROPEAN DATA-LINK ASSESSMENT, 1987

European Data-Link Assessment

Terry Bisbee William Bostwick Robert McGrath
William Miner Steve Napear Walter Sadowski
William Wing

April 16, 1987

Contents

I. Introduction

This report summarizes the operation of the test data communication link between selected U. S. and European fusion laboratories. This test was successful in that it provided a database on which an informal technical decision can be made. This report contains recommendations which are based on the test results and the committee's perception of long-term advances in networking technology and its impact in the Fusion program. Options for the future of the European link are presented in this report.

II. Scope

The material presented herein is based on a test conducted over a twelve month period during which U. S. and European laboratories could communicate via an electronic data link. Both European and U. S. users were selected from among the participants in international exchanges. Our conclusions and recommendations are based on the operating experiences of these forty users, additional reports from the U. S. - Japanese link users and non-O. F. E. network expertise. To ensure a broad base of technical expertise, the review committee included members from outside the Fusion community.

III. Background

A. MFE Program Plan for International Cooperation

(Walter Sadowski)

B. Implementing Agreements in Place

(Bill Bostwick)

C. Technological Developments in Networking

1. Introduction

The technology to support digital global networking is already in place. Its foundations lie with a combination of developments in a variety of supporting disciplines including satellite technology, fiber optics communications, opto-electrical transceiver technology, digital IC fabrication, microcomputer technology, and the standardization of many layers of communications and computer interface protocols. A prediction of future capacities and capabilities for global networking is really a game of "when" rather than "if". Rational observers would agree that global networking of digital data has become a necessity in today's environment of global commerce.

Technological forecasting has become a necessity for most firms that depend on "high technology" products and/or services. Organizations as diverse as firms ranging from SONY and PHILLIPS, who produce commercial retail products such as optical disc recording devices, to the U. S. government's Strategic Defense Initiative Program Office, expend major resources to track and predict technology. However, technological forecasting for global networking, albeit interesting, is not fully relevant to this discussion. The technological base necessary to support large digital global networking is currently in place and is just beginning to be utilized. The near term metamorphoses, over the next five years, will be economically driven, rather than technologically driven.

The key to successfully predicting where digital global networking will be in five years lies with tracking the developing standards. Knowledge of the standardization of communications protocols, possibly including actively participating in the development of these standards, will yield the clearest insights.

The organization that first plans how it will use this capability and then a priori develops and adheres to its own set of user data definitions and communication protocols can avoid much of the cost of ad hoc standards. Encumbrances from the past are the most costly. Poor or no planning can be the costliest of legacies. The following paragraphs present a historical synopsis of networking leading up to the current state of digital global communications.

2. Background

There has always been a "layering" of simple architectural features when implementing computer systems. During the 1970's, a major transition took place in the practical way computer systems were fabricated. Their internal components consisted more and more of production components, rather than components specially tailored for the application. Although the primary rationale for use of production components was economic, a secondary effect, perhaps even more profound than that of low cost computing began to cause a rethinking in the way systems are viewed.

This secondary effort first appeared in simple serial communications, RS232 class, because some of the first large scale integrated circuits were UARTs (Universal Asynchronous Receiver Transmitters). In order to lower the cost of these components through large scale production, communications protocol standardization was required. The key to success was directly proportional to the degree of standardization.

A similar phenomenon took place with the advent of standard buses. At first these public, or "opened" buses were outside the mainstream of the industry. They were primarily found in the newly evolving micro-computer industry. However, it soon became apparent that large sub-markets could be effectively captured if a integrated circuit manufacturer's own bus standard became a de facto industry standard.

Almost over night the industry proclaimed the benefits of standardization. Committees were formed to define the standards for everything from IC chip pin layouts to Local Area Network protocols. And unlike the software standardization that took place a decade before, (e.g. ANSI standard Fortran, CODASYL, COBOL, et al.), these hardware standards needed to be completely defined and completely adhered to. The economic significance of a customer being able to purchase a third party board level product, plugging it into his standardized bus computer without any trouble, drove the level of standardization compliance to this more than satisfactory state.

Today the computer industry has standards for IC chip electrical and physical form, board sizes, bus mechanical, bus electrical and bus protocols, communications protocols for inter-computer systems located in proximity, Local Area Network Protocols, and communications for inter-systems geographically dispersed, Wide Area Networks and Global Networks.

3. Current Technology Assessment

In the late 1980's the computer and data communications industry has both the necessary technological base and the standardization of communications protocols to provide full global networking on a very large scale. Multi tens of billions of dollars of new marketplaces have been envisioned. The only impedance has been the speed at which the user community can absorb these capabilities.

Progress has not been slow. The benefits in cost savings and productivity gains, coupled with the rise in global economic competition is forcing the user communities to rapidly integrate these new capabilities into every facet of their activities.

Wide Area Nets (WANs), span ranges from just beyond the few kilometer limits of Local Area Networks to full global communications. WANs are implemented with two technologies, satellite and terrestrial. As of this date, (1987), there are almost 1000 satellite transponders. Each transponder operates in the Gigahertz range, (C band at 4-6GHz and Ku band at 12 to 14GHz), potentially providing very high bandwidth communications capability. T1 class service is provided in units of 1.544Mbps, with subdivided service at 56Kbps, and again subdivided into 9.6Kbps. T2 and T3 class service provides far greater bandwidth. The obvious maximum digital rates available on the C band are in the Giga-bits per second range. However, there are currently few applications requiring greater than T1 service since rates above T1 begin to approach computer internal bus speeds.

Terrestrial communications have recently undergone a revolution in capacity due to the introduction of fiber optics communications. A single strand of fiber optics cable can handle the entire telephone communications requirements of the Northeast corridor! The limitations to realizing the tremendous bandwidth of fiber optics communications lies primarily in the electronics which convert the light signals to and from electronic signals.

For global networking, the advantages of terrestrial communications is physical, .i.e., the elimination of the up and down link synchronous satellite delay of a quarter second. The primary disadvantage currently is the physical distance most sites are from the fiber optic trunk lines. Just as dedicated satellite dishes are employed, dedicated server lines are considered when large data transmission requirements exist.

4. Standards

Networking standards provide the user community with a cost effective and practical method for interfacing between network components at a variety of levels. Standards do lag several years behind technology for several reasons. The most important and relevant is that we don't really know what we need nor how best to implement a set of functions until we try them out. This trial and error approach is really not due to a lack of planning or analysis. It recognizes the fundamental limitations of humans in our ability to understand how best to apply new technologies and techniques in order to improve our environment.

The current state of network standards includes a well defined set of official and semi-official protocols for inter-computer networking on the local level, and for packet switching on a more global level. The current activity in standards for local networking involve defining the higher level protocols: the session and presentation layers of the ISO standard in various forms. Heavy emphasis is being placed on graphics presentations and in inter-operating system command translators.

As different "experiments" are tried in the marketplace, one or more approaches will evolve as the standards. When it becomes apparent that one or more of these standards will be around for at least several years, the standards are first set into ROM, and then dedicated integrated circuits are designed to implement these standards. This usually transpires in conjunction with formal standardization through either the IEEE or ANSI.

Global packet switching protocols have been around for a relatively long time, but are based on a time sharing mode of usage. This implies their data transfer rates are proportionally slow and response is ultimately based on very slow, serial communications rates.

IV. Status

A. History of the Test

In August of 1984, at the urging of Dr. Walter Sadowski of the Department of Energy's Office of Fusion Energy, a dialog was started with three European fusion laboratories on the desirability and possibility of establishing a test version of an electronic computer-to-computer link between the U. S. and European Fusion programs. Logistical and technical details were discussed during a visit by J. Leighton (of Lawrence Livermore National Laboratory), L. Mann (of Los Alamos National Laboratory), W. Wing (of Oak Ridge National Laboratory), and W. Sadowski to the European fusion Laboratories. The logistical problems could potentially have been considerable since the European telecommunications carriers are all government owned. They explicitly preclude users linking computers over their systems unless they use the approved X.25 protocol and licensed hardware. Since the European labs (IPP-Garching, KFA-Jüelich, and Culham) all had VAX computers available, an expedient solution to this problem was available based on using off-the-shelf commercial (licensed) VAX software and hardware. DEC offers a package for VAX's that provides the packet assembly-disassembly (PAD) function necessary to use the X.25 protocol. By installing this software on VAX's at each site, and using one of the approved (licensed) hardware interfaces available from DEC, direct connection could be made to the national X.25 networks. Through these X.25 networks, connection could be made to the international X.75 trunks that would carry the test.

This procedure was agreed to, and work began toward implementation. In February of 1985 funds were added to ORNL's financial plan to allow procurement of hardware and software and to allow contracting with TYMNET to provide gateway service to the international X.75 network. In July of 1985 hardware was delivered to ORNL and in August the PSI software was delivered. Procurement proceeded even more slowly in Europe. The funds at ORNL were carried over from FY-'85 into FY-'86. In November of 1985 the first electronic contact was made with Europe, and in December the first user traffic was carried. A limited number of users (roughly 40) were solicited from among participants in various exchanges and extended laboratory working visits. By May of 1986 the traffic generated by these users had stabilized at roughly 20 hours of connect time per month. In December of 1986, a formal review of the test link was initiated by the Office of Fusion Energy. This report is the product of that review.

B. Present Capabilities

The present capabilities of the link reflect its status as a test installation with a limited user community and limited throughput. It should be noted that the hardware capabilities of the link are largely independent of the software capabilities. Hardware support for link traffic is provided by two terminal ports on a VAX 780. These ports are connected through leased lines at 4800 baud to the closest TYMNET gateway. This in turn provides the actual connection to the X.75 international trunk. Thus, although the total aggregate throughput is 9600 baud, the maximum throughput seen by a single user is 4800 baud. If multiple users happen to be simultaneously sharing a port, their traffic is multiplexed and the 4800 baud is shared. The VAX 780 is a node (known as ATF) on the MFE network and thus serves as a bridge between the MFE network and the X.75 trunk.

The software support for the link is provided by DEC's PSI (Packetnet System Interface) which is layered under DECnet and provides the packet assembly-disassembly (PAD) functions necessary to connect to an X.25 network. It transparently supports all DECnet functionality, or any subset of that functionality desired. For the purposes of this test, PSI was allowed to provide virtual terminal (SET HOST), and electronic mail functions. Since DEC's electronic mail provides the ability to send text files as well as messages, it was decided not to turn on the standard DECnet remote file access features. This decision (made at the urging of computer oversight staff in the Office of Fusion Energy) guaranteed that no files or data could be compromised by remote access from Europe. The software capabilities of the link are defined by the intersection of the set of capabilities of the MFE network software and the PSI/DECnet software. A user at a remote MFE site who desires access to one of the European labs must first login to ATF via the MFE network. Having done this, the user can read mail, send mail, or SET HOST to one of the VAX's at the participating sites in

Europe. Obviously, the user must have an account established on ATF and at any VAX to which virtual terminal access is desired. Mail however, can be sent or received without a user needing to have an account in Europe. It would be possible to install a version of the MFE mail system that could automatically forward MFE mail to PSI mail and vice-versa. This has deliberately not been done however, because doing so would completely prevent any sort of tracking of sources and volumes of mail forwarded from remote sites. We have instead set up the PSI software to allow very close control of network use. The PSI software checks a "rights" database in the VAX before allowing the user to perform either a PSI-MAIL or a PSI-SET HOST operation. These "rights" allow controlled access to the link, control the remote sites to which individuals can get access, and could (if desired) monitor attempts by unauthorized people to gain access to the PSI-link. In addition, separate accounting is run on all PSI sessions. The PSI accounting files log each incoming and outgoing mail message, and each incoming and outgoing virtual terminal session. The total connect time, and total transmitted and received byte- and segment-counts are recorded for each session.

C. Examples of Use

After one year of operation, users of the European Gateway were asked to provide their impressions of its usefulness and effectiveness. In all cases researchers stated the productivity of their international efforts were greatly enhanced by access to the link. Most remarkable though, is the variety of activities of which this link has become a central part. Some of these activities are briefly described below. More detailed information on current usage of the link can be found in Appendix A.

As one would expect, theoretical researchers and computational modellers traveling between U. S. and European laboratories find the link extremely useful since it provides them with easy access to their files at the home institution. This access enhances the productivity of these traveling researchers in several ways. Certainly, it greatly simplifies the advanced planning for foreign travel. One no longer has to guess what codes and files might be needed, transfer these to tape, and re-establish them on another and often different computer. Further, the link keeps the traveling researcher in contact with his colleagues at the home institution. This tends to establish interaction between research groups rather than between individuals. At the conclusion of a foreign visit, the link ensures continued collaboration by providing convenient communication and easy access to new capabilities available or developed at the foreign facility.

Experimentalist traveling between U. S. and European fusion research laboratories has found this link to be an essential part of their operations. Once a diagnostic or component has been installed on an experiment at a foreign laboratory, the developer of the equipment can use the link to monitor its operation and perform data reduction from his home site. For example, researchers at UCLA, ORNL and Sandia regularly log on to the VAX computer at KFA Jüelich to reduce or analyze data from ALT-I and its associated diagnostics. In some instances the link is being used to provide data on a real time basis. For example, Dr. Haas from IPP Garching has provided a fast ion gauge for use in the ergotic limiter experiments on TEXT. For critical experiments, Haas monitors the data acquisition for each tokamak shot and provides direct feedback that assists in specifying conditions for the next shot. A similar example exists within the IEA sponsored Large Coil Task where Dr. A. Ulbricht from Karlsruhe and his colleagues monitor, in real time, experiments at ORNL. This allows for review and inspection of the results by experts at Karlsruhe.

The European Gateway Communications Link will serve as a central element of other U. S.-European collaborative efforts involving large numbers of scientists. As an example consider the Advanced Limiter Test II, which involves SNL, UCLA and ORNL from the U. S., IPP Nagoya and MHI from Japan and KFA Jüelich, FRG. The U. S. has contributed $ x million (Oktay can provide this number) to the development effort on ALT-II over the past three years and the effort has involved nearly a hundred different U. S. scientists. Operation of this experiment begins in the spring of 1987. The U. S. plans to provide at least four researchers stationed at TEXTOR throughout the first year of operation with as many as 20 additional U. S. personnel traveling to KFA for shorter visits to participate in the experiments. The European Gateway provides all of these scientists visiting KFA with information about the tokamak operation prior to their visit and easy access to additional data after their return home. Communication between scientists at the experiment and their co-workers at the home institution is even more essential for an experiment of this magnitude since it provides all the participants in the design effort with direct access the experimental data. Thus the large modeling efforts in plasma confinement, plasma edge dynamics, plasma materials interaction, limiter pumping and neutral transport, and many other areas which contributed to the design, can be benchmarked against the experimental data and can be used to contribute to the analysis of the operation.

The U. S. is also conducting an even larger collaborative effort on Tore Supra, with experimental operation scheduled to begin near the end of 1987. The U. S. is contributing to this program in the areas of pellet injectors, pumped limiters, and rf antennas. At present, the European Gateway Communications link does not include CEA Cadarache. This is extremely unfortunate, since the experience on ALT-II has been that good communication, including access to each others analysis and design codes, minimizes confusion, mistakes and redundant work. This greatly assists in avoiding schedule delays and directly impacts costs.

Although the use of experimental data by visiting scientists is extremely important, it will play an even more important role in the O. F. E. Theory Program as a way to test theoretical models against experimental results. The need to test and evaluate theoretical models has become imperative due to the increasingly important role that transport simulation has played in the predictive modelling required for advanced tokamak design. As machine construction costs increase, improved confidence in the results obtained from these simulations is required.

To this end, a database of experimental data to which the theoretical models can be compared has been established. The MFE database of experimental results is patterned after the database currently in use at the Fusion Research Center at The University of Texas at Austin to maintain data obtained from the TEXT experiment. The MFE database will provide the MFE community with complete sets of experimental data from a variety of machines and discharge conditions. For the first time a common set of experimental data sets will be available to the entire MFE community via the MFE computer network. Although the primary application addressed here is the comparison between theoretical models and experimental results, there are also many additional uses for such a database.

Maintenance of the database requires retrieving new data from available sources to broaden it scope. The data link between the U. S. and European fusion laboratories will supply additional sources of experimental data to the database. Not only can data be retrieved from those experiments mentioned above, but also from JET at Culham and both ASDEX and WENDELSTEIN at Garching, as they are already accessible via the European Gateway. One such data set has already been shipped across the network and stored in the database. The electronic transfer of data into the MFE database will eliminate time consuming and possibly inaccurate digitizing of graphical data.

The availability of such a database throughout the MFE community, not only in the U. S. but also in Europe via the European data link, should also stimulate new interests in the area of predictive transport modelling from scientists who previously did not have such a database of experimental results available. This should not only stimulate scientists, new to this area of research, but also should hopefully create new collaborative efforts between research sites throughout the community. The data link and the MFENET provide a path along which manuscripts, ideas and data can easily be transferred and shared among it users.

The documented operation of the existing European Gateway Communications link shows that it is already enhancing the productivity of U. S. collaborative programs with European laboratories. Cooperative efforts such as those on TEXTOR and Tore Supra are just beginning to reach their peak level of activity and this link, if available, will play a central role in these programs. Given the large investment of U. S. funds in cooperative efforts with the European laboratories and the extremely low costs of the link, there appears to be little question of its cost effectiveness.

V. Requirements

A. Long Term ER/OFE Network Requirements

The Office of Energy Research (ER) requirement for data communications has grown over the years to serve various needs. The most significant milestone in this regard was the creation by the Fusion Energy program (OFE) within ER of the National MFE Network during FY 1975. The MFE Network was initially implemented to provide access to centralized supercomputer resources for fusion energy researchers throughout the U.S. However, the MFE Network has evolved to tie together computers of differing capabilities and to support collaborations and applications for design engineering and experimental facilities. The scope of the MFE Network was broadened to include the non-fusion ER programs during FY 1984 with the expansion of the ER supercomputer access program. In concert with expansion of the supercomputer program, a similar broadening of capability requirements across the larger ER community has emerged. Based on these requirements, ER has implemented a plan to interconnect all ER facilities through an integrated computer networking infrastructure referred to as ESNET. ER's goals for the ESNET concept are as follows:

- to support interdisciplinary research collaborations and communications,
- to provide flexibility and interoperability for ER multi-vendor computer operations,
- to provide access to ER research facilities (including supercomputers, experimental facilities, etc.),
- to evolve toward interagency computer network standards,
- to establish interconnectivity with computer networks supported by other government agencies, and
- to avoid redundant computer network support costs.

ER has also recognized the need and requirements for European and Japanese data links to support ER involvement in several recognized international research collaborations and is currently experiencing a phenomenal growth of these requirements as large scientific collaborations have become more international in scope. At present, ER scientists collaborate on experiments and projects at research facilities in Switzerland, Italy, and Japan as well as in Germany via the European test data link described in this report. Of the existing collaborations most of the requirements have been driven by the Fusion and High Energy and Nuclear Physics programs. However, a significant increase in both existing programs as well as new requirements in Basic Energy Sciences and Health and Environmental Research are taking place in the fiscal years 1987 and 1988 time frame.

B. Assessment of Long Term European-Link Needed Capability (MFE Global Networking)

The MFE funding profile has seen a severe downturn during the last several years. This loss in funding has been exacerbated by the steadily rising cost of each new experiment. An obvious aid to this problem lies with cost sharing of new experiments among several nations. Cost sharing naturally implies sharing of the experiment between scientists of these several nations, with the subsequent reduction of experiments within the direct control of DOE. The less experiments, the fewer direct access available to our scientists and researchers without heavy overseas travel.

This scenario aside, there are several major extra-continental experiments that we exchange our researchers with theirs. Since we still have the dominance in computational resources, available computer resources are more readily available here, especially for any detailed second generation analysis. A fundamental question is whether it is cost effective, that is, whether more or better science can be accomplished by having the experiment physically separate from the long term (or even short term) analysis tools.

This appears to be a difficult question to answer. Much of the value of the gathered data is so dependent on the "non recorded" circumstances of the series of shots that some claim it is difficult to understand the data without "having been there." Yet we save and archive as much of this data as technologically possible, spending significant sums on this storage capability. Images of large vaults within the fission waste sites in the bowls of the Nevada desert for DOE fusion databases come to mind.

Yet at least some of this data, or its reduced form, does have meaning and the proper presentation and manipulation of these data sets potentially can yield great insights into the fusion phenomenology. With the proper set of analysis tools, the scientist and researcher should be able to glean insights by analysis of either single experiment data sets or of multi-experiments data sets.

The computer industry is on the verge of a third generation of workstations, i.e., tools for individuals within a given application to substantially increase their productivity. These workstations will range in cost from the least expensive personal computers up to the Sun workstation class machine. They will include a much higher resolution display with the electronics (graphics processors, wide bandwidth DAC's etc.), to provide a significantly greater level of functionality than current 2nd generation workstations. Their "main" processors will be range in performance between 5 and 10 MIPS, (about a VAX 8700 class capability), and have a local area network capability to transfer large data sets almost instantaneously.

Regardless of the OFE funding profile, (within bounds!), DOE fusion researchers will find themselves working at these workstations. They will find themselves traveling (overseas) to gain the necessary access to fusion experiments. Data gathered will be sent back with the researcher, and most likely will be transmitted over some global data network.

The choice is not whether fusion research will work this way but rather how can OFE best guide the evolution to minimize the cost and maximize the gain. OFE should be working so that the minimum amount of wasted effort is expended. In five years we should find ourselves not with many camps of individuals, each working using his or her set of software, incapable of sharing data and software, but with an optional set of OFE standards, software products and tools. OFE should be providing the guidance and leadership to develop these standards.

C. Near- and Long-term MFE requirements

Near-term and long-term requirements for a permanent network can be broken into two categories. The first category includes requirements for security (or integrity), reliability, and limitations on access. These requirements are presumably rather static and not a function of current technology or implementation. The second category includes the desired functional capabilities the network would have for supporting work. This second category should be regarded as a dynamic goal rather than a fixed set of requirements. At any given instant of time, people's requirements will be set by the working patterns they have developed. These, in turn, are shaped by available technology and will evolve over time.

The first category (security-related) requirements can reasonably be modeled by any modern operating system. Like an operating system, the network must provide:

1. Access Restrictions

Local (at each site) management policies must be able to place restrictions on the ability of local users to place out going calls. These restrictions should be controllable on a user-by-user bases and should be able to restrict the locations a user can call, the volume of traffic a user can generate, and the total time a user can be connected. Similarly, each site must be able to control remote access to it. A site must be able to control which external sites can successfully connect to it, what external users can successfully log in, and finally, what privileges such external users are permitted (again on a user-by-user basis). The privileges granted external users should form a clear hierarchy from none (send in electronic mail only), up through being able to connect a virtual terminal to an account with limited quotas or privileges, connect to a normal user account, and finally be able to exercise remote direct file access at the record level. This highest level of access would permit remote direct access to database facilities.

2. Data Integrity

No user should be able (either accidentally or maliciously) to be able to corrupt another user's network traffic. It should go without saying that the network itself should have a bit-error rate at least as good as typical magnetic media, and should detect and correct errors in a transparent fashion.

3. Reliability

The network reliability (availability) should be at least as good as the joint availability of combinations of nodes on it.

The desired functional characteristics of a network will evolve as people develop work patterns based on the evolution of local systems. Since at the present time, the network which MFE people are most accustomed to (and therefore use as a paradigm) is an Ethernet-based DECnet. This should serve as the definition of the functional goal for a present-day network. This goal can be defined in terms of several functional features or measurable quantities.

4. Communication Speed

For both interactive and batch work on remote computers, users require adequate communication speeds. Experience on local computer networks manipulating text indicates that communication speeds equivalent to 9600 baud or higher is required to effectively utilize both machine and human resources. For support of complex interactive graphics or transmission of bit-mapped images higher speeds (up to the 1.2 Mbyte Ethernet limit) are required

5. Electronic Mail Services

Electronic mail provides an effective means of communicating with other scientists or groups of scientists in an expeditious manner. Electronic mail greatly facilitates collaborative efforts. An electronic mail service should provide automatic logging of mail received, automatic sender notification of mail receipt, and automatic forwarding. Mail forwarding should be capable of reaching all nodes on the network regardless of their connection path. Finally, an electronic address book should be maintained on-line. This will make it possible to send electronic mail without prior knowledge of the intended recipient's account name or node.

6. Terminal/Workstation Access

Terminal/workstation access must be provided in a transparent manner. Network latency must not impede terminal/workstation access to remote computers. This places stringent constraints on either the latency time of the network (making satellites unacceptable) or the level of cooperation between the workstation and the host (e.g., the workstation must make intelligent decisions about when to turn on or off local-echo). A defined standard applications interface (for example X-windows) must be supported in a transparent manner. As available bandwidth allows better support for local workstation processing, a standard remote file access protocol must be supported (for example Sun's Network File System, or DEC's Record Management Services).

7. File Transfer

The file transfer protocols should support all file types.

Both prioritized and interleaved file transfer should be available so that long files do not unduly impede the flow of short files. It might prove cost effective to have two network channels available, a high bandwidth (but slow response) link which could be used for file traffic and a fast response channel for interactive traffic. Users should be provided with the means to schedule file transfers for off-hours transmission (and encouraged to do so!).

8. Remote Printing/Plotting

The user should be able to direct files to any "supported" printer/plotter on any node on the network for printing/plotting (consistent with the access constraints mentioned above). This may require the adoption of a standard metafile format.

9. Interprocess Communication

Support for interprocess communication is undoubtedly necessary for implementing other requirements, e.g., for the efficient implementation of full screen editing, workstation support, etc. While the development of these capabilities is beyond the scope of the network per se, the design of the network must not preclude nor impede these applications. This capability is also likely to be important for distributed CAD/CAM applications as well as the research activities envisioned here.

10. Remote Job Entry

Job entry should be supported on any computer system on the network for which the user has access. Similarly, remote support for "SCRIPTS", "EXECS", "PROCS" or similar command streams should be provided.

11. Network File System

The network file system should allow access (consistent with the access requirements above) to the contents of remote files without file transfer. This access should include read/write capabilities at the record level.

12. Document Transfer

The network should permit the transfer of scientific documents. This is more than the transmission of text files. A standard Document Interchange Format (DIF) should be used to transfer formatted documents with merged text, graphics, images, and data.

VI. Options and Alternatives

A. Shutdown

The link could be terminated with minimal effort. The present contract with TYMNET contains a penalty clause for cancellation with less than three months advance notice by either party. The cost of this penalty is roughly the dollar value of retaining the leased line(s) on idle status for the same three month period (about $9,000).

B. Maintain

The link could be maintained at its present level, both in terms of number of connections and in terms of throughput. At this level the link has served to connect the U. S. program to three European sites for over a year and a half. It has served the communications needs of over forty traveler-users and has proven itself to have had a benefit-to-cost ratio significantly greater than one. Thus, it could continue in this mode (with roughly the same size user community) for roughly the same cost ($75,000/year).

C. Upgrade

The link could be upgraded, both in terms of the European laboratories to which connections were supported, and in terms of absolute throughput. There are a number of additional European laboratories which it would be desirable to include in the "network". They include Cadarach, where TORE-SUPRA is being constructed and where there will be a strong U. S. presence; Madrid, where the Junta for Nuclear Energy (JEN) is constructing TJ-II, a flexible Helliac; and Padua, where the Italian Institute for Ionized Gases is constructing a reversed field pinch. The maximum throughput of the present link is 4800 baud. This is marginally acceptable for text transmission and remote terminal access to text files. It is unacceptable for graphics and data transmission. Since the throughput the present MFE network offers to a single terminal user is roughly 4800 baud, supporting a gateway connection with bandwidth significantly greater than this was not cost effective for the purposes of a test-link. The test rarely involved more than one user at a time. However, the advent of MFE-Net

II should provide significantly greater terminal bandwidth than 4800 baud, and if more users will be using this bandwidth to access more laboratories and be attempting more work (particularly graphical data analysis), then the gateway bandwidth would also have to be increased.

VII. Recommendations

A. The committee recommends selection of option three, continuation and upgrade of the link capabilities.

B. A standing committee be formed to oversee the development of a set of experiments to gain insights into the optimal working environment for experimental research using workstation technology. These experiments have a direct bearing on the requirements for global networking, since the requirements for the manipulation and presentation of data drives the requirements for storage and exchange of data.

C. Several of the DOE laboratories should be designated as experimental centers, with several concurrent experiments ongoing simultaneously.

D. A semi-annual review and presentation of the results and a decision for the subsequent experimentation should take place.

E. From these experiments, as it becomes obvious that a set of standards will be required, and as the experiments yield insights into the appropriate (optimal) form of the standards, a standards should be formed and chartered to evolve and maintain the particular standard.

F. Standards for data storage format, including trying to identify those elusive "have to have been there" characteristics should be immediately commenced. Archived data that cannot be understood for lack of characteristics is most probably wasted.

Reference 9

ESNET X.25 MIGRATION PLAN, 1987

ESNET X.25 Migration Plan
08 APR 1987

PROJECT DEFINITION

This is an effort to provide an interim, X.25 based, backbone facility prior to the installation of such capability on the ESnet for the HEP community.

INTENTION

This backbone facility is considered to be part of the ESnet internet, although it may (or may not) be disjoint from the IP based ESnet backbone facility.

This backbone facility is NOT considered to be an alternate or potential alternate to the integrated ESnet concept. It is considered an interim solution and will be integrated in a manner that is consistent with the goals of ESnet.

Recognition of the interim nature of this effort implies the following:

1. Operational status must be achieved in the very near term.
2. Costs should be consistent with the short term nature of the project.

TECHNICAL DESCRIPTION

The facilities to be provided will consist of 56K bps dedicated links from LBL-FNAL, LBL-MFECC, FNAL-BNL, BNL-MIT. X.25 switches, associated hardware, and management facilities will be included.

LEP3NET will be interconnected at MIT.

The CERN satellite link will be interconnected at FNAL.

The link to KEK, Japan will be interconnected at LBL.

The FSU-FNAL link and switching facilities at FSU are not directly included in this effort, but coordination will be required to ensure compatibility.

RESPONSIBILITIES

A. Requirements: The Requirements team will be responsible for specifying the HEP community requirements including:
    1. Identify technical issues
    2. Specify technical requirements and specifications.
    3. Identify possible alternative approaches.
    4. Recommend switch configurations for each site.
    5. Review the procurement strategy and specification (if any) for comment.
    6. Provide liaison support for installation and initial operation.
    7. Develop a schedule and order of installation.

B. Procurement: The Procurement team will be responsible for the procurement activity. This team will conduct the procurement including:

1. Select the procurement strategy consistent with technical requirements and budgetary, time, and interim nature constraints.
2. Develop a procurement specification if required.
3. Develop a selection criteria if required.
4. Perform equipment testing if required.
5. Perform vendor selection.

C. Installation and Operation: The ESnet Control and Maintenance team from the NMFECC will be responsible for:

1. Ordering and installation of communications links.
2. Coordinating hardware installation
3. On-going operational management of the backbone facility.

PERSONNEL

A. Requirements team:

1. LBL personnel to lead: Marv Atchley
2. CAL TECH representative: Harvey Newman
3. CERN technical respresentative: Bob O'Brien
4. FNAL representative: Greg Chartrand
5. MFECC technical personnel: Rick Schnetz/Tony Hain
6. SLAC representative: Les Cottrell
7. BNL representative: Graham Campbell

B. Procurement Team:

1. MFECC technical personnel to lead: Tony Hain
2. CERN technical respresentative: Bob O'Brien
3. MFECC technical personnel: Rick Schnetz
4. LBL personnel: Marv Atchley
5. HEP personnel: (to be determined)

C. Installation and Operation: The ESnet Control and Maintenance staff involved in this effort will consist of:

1. Tony Hain (ESnet Network Control and Maintenance Project Leader)
2. Rick Schnetz.

PROCEDURE (some of the following may happen in parallel)

A. The Requirements team should form and conduct meetings to formulate requirements and/or specific specification items, possible alternatives, site configurations, and schedule.
B. The Procurement team will meet with potential vendors to overview costs and technical capabilities.
C. The Procurement team will define switch testing procedures, IF required, and conduct testing of equipment.

E. The Procurement team will develop a procurement plan, accounting for technical requirements, schedule, configuration, procurement, and budget constraints. The plan will be reviewed by the Requirements team for comment.

F. The Procurement team will conduct the procurement.

G. The MFECC will be responsible for installation of the system and bringing it to a fully operational state.

IMPLEMENTATION

The NMFECC will be responsible for the acquisition and installation of the switches, links, and associated hardware. The NMFECC will bring the X.25 backbone to a fully operational state. Support and liasion with site technical personnel for both installation and ongoing operation will be required.

Health & Environmental Research

Walter Goad - Los Alamos National Laboratory (LANL)

Scientific Computing

Sandy Merola - Lawrence Berkeley Laboratory (LBL)

ESNET Implementation

James Leighton - Lawrence Livermore National Laboratory (LLNL)

The first meeting of the ESNET Steering Committee was held on December 3, 1986. Dr. James Decker, Deputy Director, Energy Research, and John Cavallini, ESNET Program Manager, of the Scientific Computing Staff welcomed the steering committee, provided background information, and led discussions concerning ESNET and its goals. The Steering Committee was asked to review the planned backbone configuration, to identify current program requirements, to set priorities for implementation, and to consider the need for any other recommendations for the ESNET. In addition, the members were asked to begin development of longer term program requirements for computer networking.

At the second meeting held on January 29 and 30, 1987, the committee agreed by adding interim 56 Kbps links connecting LBL-Fermi-Brookhaven-MIT should be incorporated as part of the ESNET backbone configuration to meet current High Energy Physics (HEP) program requirements, and to establish an early phase X.25 transport mechanism as part of ESNET. This interim solution was an agreed first step in the migration of HEPNET into ESNET pending the provision of these services via announced IP/X.25 gateways. A working group was formed consisting of Leighton, Brandenburg, and Loken to work on a plan for implementing the interim X.25 links. In addition, the committee agreed on a major high speed (56 Kbps) node at CEBAF.

The committee also investigated international networking requirements, especially links from LBL to Japan, from Fermi to CERN, and the connectivity requirements for ER programs to Germany. The steering committee agreed to pursue the goal of single point satellite links to foreign countries that would serve all programs.

A second working group consisting of Dunning, Loken, and Greenwald were asked to develop written performance characteristics and functionality requirements for ESNET. [6]

At the third meeting held on March 26 and 27, 1987, Jim Leighton presented an implementation plan to support X.25 applications on an interim basis, and to provide a migration path for the large DECNET requirement of the HEP program until the new network hardware and software is deployed and is able to support this requirement. Jim also presented an implementation plan to proceed with previously planned 56 Kbps links from FSU to PPPL, ORNL and to the University of Texas. The steering committee, after discussion, endorsed these implementation plans. The committee further agreed that anticipated traffic patterns warranted upgrading the CEBAF link to 56 Kbps and to redirect the CEBAF connection to PPPL rather than to ORNL. The steering committee endorsed recommendations citing the need to define a network infrastructure to participate in the ongoing management and operation of ESNET and to provide better documentation on ESNET and procedures for accessing ESNET in the DOE grants and contracts process.

At the fourth meeting held on June 18, 1987, the steering committee reviewed the status of the interim X.25 56 Kbps backbone and the 64 Kbps satellite link to CERN, currently planned for initial operation in November of 1987. Also discussed was the status of progress reached during a June trip to Europe to initiate planning for ESNET requirements abroad. The ESC also agreed on final editing and changes to the ESNET Functional Requirements document to be incorporated into the ESNET Program Plan.

The primary purpose of this meeting was to review and comment on the ESNET Program Plan. The attending members unanimously supported changes to provide a stronger endorsement for the funding of ESNET. Additionally, the members agreed that a statement be added to the plan acknowledging the ESC members contribution and endorsement.

V. ESNET Infrastructure

There are many components which define the ESNET infrastructure and these must be put into place to provide the framework for planning, managing, and operating such a large computer network. The need for this infrastructure cuts across ER program areas while requiring the participation of these same ER program areas.

People are perhaps the most important component of the ESNET infrastructure. A nationwide computer network to support multiple research programs will require the Scientific Computing Staff to coordinate requirements, to develop plans and budgets, and to manage both the operational aspects of the ESNET and also the research and development program necessary to evolve the ESNET to meet future requirements and goals. It also requires the participation of ER headquarters program organizations, i.e., HENP, BES, FE, AMS, and HER, to identify and set priorities for ER program requirements for computer networking and the participation of the ER Automated Office Support Systems (AOSS) program manager to identify the research computer networking component of the ER headquarters AOSS requirement. The ESNET infrastructure includes the computer network group at the NMFECC who are responsible for research, development, planning, and implementation and operational management of new capabilities or facilities, for coordinating installations, for gatewaying to other agencies' networks, and all aspects of ESNET operations. However, this centralized operation must be complemented by personnel at major network nodes who are responsible for applications level network support, for local and campus area network connectivity at their site, for selected gateway support to other communities of interest, and for local consulting and assistance.

The ESNET also requires a standing committee of users representing the ER program areas to provide direction and to coordinate requirements as described for the ESNET Steering Committee described in section IV.

Facilities form another component of the ESNET infrastructure. These facilities include hardware and software systems, data communications media, gateways to other networks, and also the architectural model which describes and structures the way in which these facilities interconnect and interoperate. In fact, the architectural model of a computer network directly affects the manner in which individual entities within the network interact. The open architectural model required by the ESNET imposes a greater need for management at individual sites, such as a university campus connected to the ESNET, to coordinate local requirements and to provide broader local connectivity/access to the ESNET and to other wide area network resources. This is necessary to avoid redundancies while ensuring that local performance objectives are met.

In the same manner, this architectural model must extend to ER Headquarters in a seamless and transparent fashion since program management for the various ER programs requires substantial involvement and participation in the research projects in the field. Hence, the ER Headquarters site should make a concerted effort to establish a local area networking capability to provide connectivity for all ER program areas and to allow effective and efficient interfacing to the ESNET. It is anticipated that a ESNET/ER Local Area Network will provide a significant management tool for ER program managers, will enhance ER information gathering and reporting capabilities, and will allow ER programs to benefit from one another's research efforts.

Another necessary component to assure widespread connectivity, is the adoption of standards for use throughout the ESNET. These standards need to cover not only the facilities, equipment, and their capabilities but also the services required for computer networking in a scientific domain. Hence, standards should be adopted for operations, procedures, applications, interfaces, user services, and much more.

The ESNET infrastructure will also require a framework which will encourage information exchange and project coordination. There should be support for regular user meetings, workshops, and other vehicles to facilitate user feedback, information dissemination, inquiries, and coordination.

A last and very essential component for the ESNET infrastructure is the establishment of and continued support for a research and development program for computer network technology especially oriented toward the ER mission. This program is necessary to ensure timely support for future scientific network application requirements.

VI. Computer Network Requirements

During the next 5 years, ER researchers will need access to wide area computer networks that are at least 1000X times more capable at the end user interface than those that are generally available to the ER community today. This capability is needed to support network functionality for remote high resolution graphics workstations, remote network file systems, distributed systems, high volume data transfer, and emerging higher speed supercomputer system interfaces. Figure VI-1, a data transmission nomogram, shows that a researcher, using a high resolution black and white workstation screen for an application which requires visualization, requires end-to-end access to data transmission capabilities in excess of T1 data circuits (1.5 Mbps).

The overall ER computer network requirement is described in five attached requirements documents for the High Energy and Nuclear Physics, the Fusion Energy, the Applied Mathematical Sciences, the Health and Environmental Research, and the Basic Energy Sciences programs, respectively. [7] Also, included in the attachments to this section are a statement of functional requirements, Desired Functional Characteristics of ESNET, prepared by a subgroup of the ESNET steering committee and a white paper, [6] prepared by D. Hall, et al at LBL. These last two papers, when analyzed in conjunction with the data transmission nomogram, show that the data communications facilities currently available to the ER user community will need to be upgraded significantly over the next five years. Because of its potential to provide cost-effective high bandwidth communications, satellite multicast technology should be given strong consideration for implementation during this period. In the longer term, full fiber data transmission technology may become available at reasonable costs.

The functional requirements identified in reference 6 are based on input from the ER research community and these functional requirements have been reviewed by the steering committee. These requirements include interprocess communication, scientific workstation support, distributed window systems, network file systems, and document transfer, which, according to reference 6, are network functions requiring much higher data communications capacities. Also included are the more traditional functions of electronic mail, terminal access, file transfer, remote printing and plotting, remote job entry, network management and control, and connectivity and interoperability with other networks. Full screen editing capabilities are vital to manage the modification and manipulation of large scientific computer codes (1 million lines or more) in an efficient and productive manner from remote locations. However, a great deal of research and development is necessary in the software implementation of these systems to improve their network performance characteristics and to broaden their functional capabilities. Also, multicast technology appears to hold great promise for functional support for scientific collaborations and for providing economical, high bandwidth capabilities to many distributed locations.

In order to provide a flexible approach to network connectivity and to interoperate across local area network or other agencies' network domains, it is necessary to support the International Standards Organization (ISO)

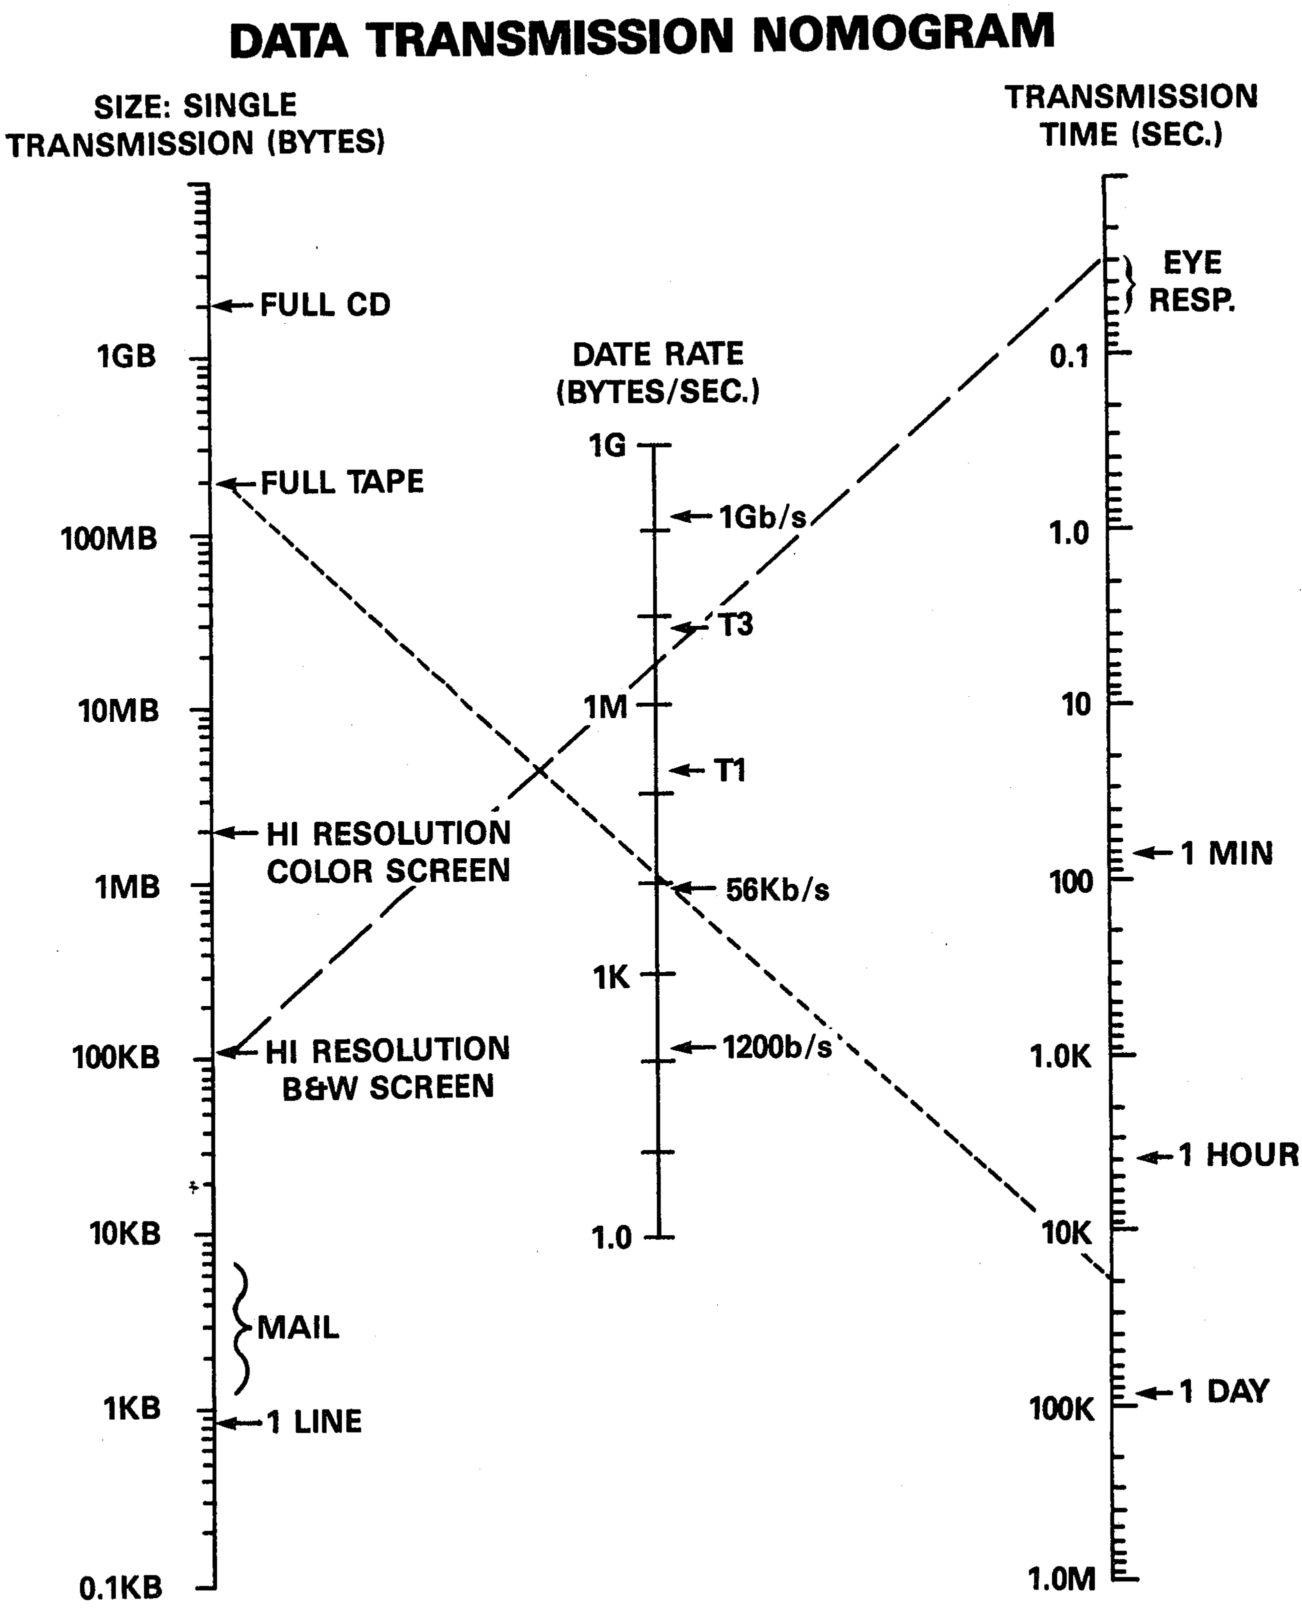
DATA TRANSMISSION NOMOGRAM
SIZE: SINGLE TRANSMISSION (BYTES)
TRANSMISSION TIME (SEC.)
DATE RATE (BYTES/SEC.)
FULL CD
1GB
FULL TAPE
100MB
10MB
HI RESOLUTION COLOR SCREEN
1MB
100KB
HI RESOLUTION B&W SCREEN
10KB
MAIL
1KB
1 LINE
0.1KB
1G
1Gb/s
T3
1M
T1
56Kb/s
1K
1200b/s
1.0
EYE RESP.
0.1
1.0
10
1 MIN
100
1.0K
1 HOUR
10K
1 DAY
100K
1.0M

Figure VI-1

network protocol suite. In the interim, as several of the attached requirements documents point out, the ER community has a substantial investment in DEC hardware and, therefore, must continue to support the use of DECNET protocols. In addition, the ER community requires continued access to supercomputer systems via the existing MFENET protocols (these were developed by the NMFECC to meet supercomputer performance objectives) and these must also be supported. Hence, the NMFECC proposal to implement an Internet Protocol based network which can accommodate multiple higher level protocol suites offers an architecture for ESNET which meets these varying requirements.

Several of the ESNET functional requirements, such as interprocess communications, workstation windows, interactive distributed graphics, etc., emphasize the need for network software research and development and for the adoption of standardized or common applications interfaces, e.g. remote procedure calls or graphics interfaces.

The requirements to provide full connectivity to the ER research community and to meet network performance objectives are particularly sensitive to cost and deployment variations. The ESNET steering committee analyzed these issues with regard to defining the ER major nodes for the ESNET backbone configuration and reached consensus for the configuration shown in Figure VI-2. Bandwidth upgrades to this configuration and the degree of deployment beyond current installations will require further analysis and will be subject to the availability of funds, the ability to implement class of service routing (satellite vs terrestrial routing) to meet performance objectives, and the extent to which other networks will lessen deployment requirements by providing network access to U.S. researchers through interoperation.

Figure VI-3 is a list of geographic locations which have been identified as the basic requirement for ESNET connectivity. This list has been compiled from the individual program requirements documents. The list of geographical locations is important for several reasons. The list totals over 100 locations which, if all were serviced with communications links of 56 Kbps, would require approximately $5,000,000 in annual funding for just this single network component. Nevertheless, as it has been noted, this capacity level is inadequate to support many needed scientific applications. It should be noted however, that most of these locations overlap existing or planned network installations for other Federal agencies, especially the National Science Foundation (NSF). Hence, it will be important to plan and to coordinate installations and service levels at these locations to minimize redundancies. For example, if two-thirds of the overlap were supported through gateways to these other agencies' computer networks, then approximately $2,000,000 less would be required for these 56 Kb circuits. When communications support for T1 circuits are considered, these costs and potential savings are multiplied. Another important statistic to note is that over forty percent (40%) of these locations already have multiple communications links that are supported by ER funds. Proper coordination and implementation through the ESNET process should help rectify this situation.

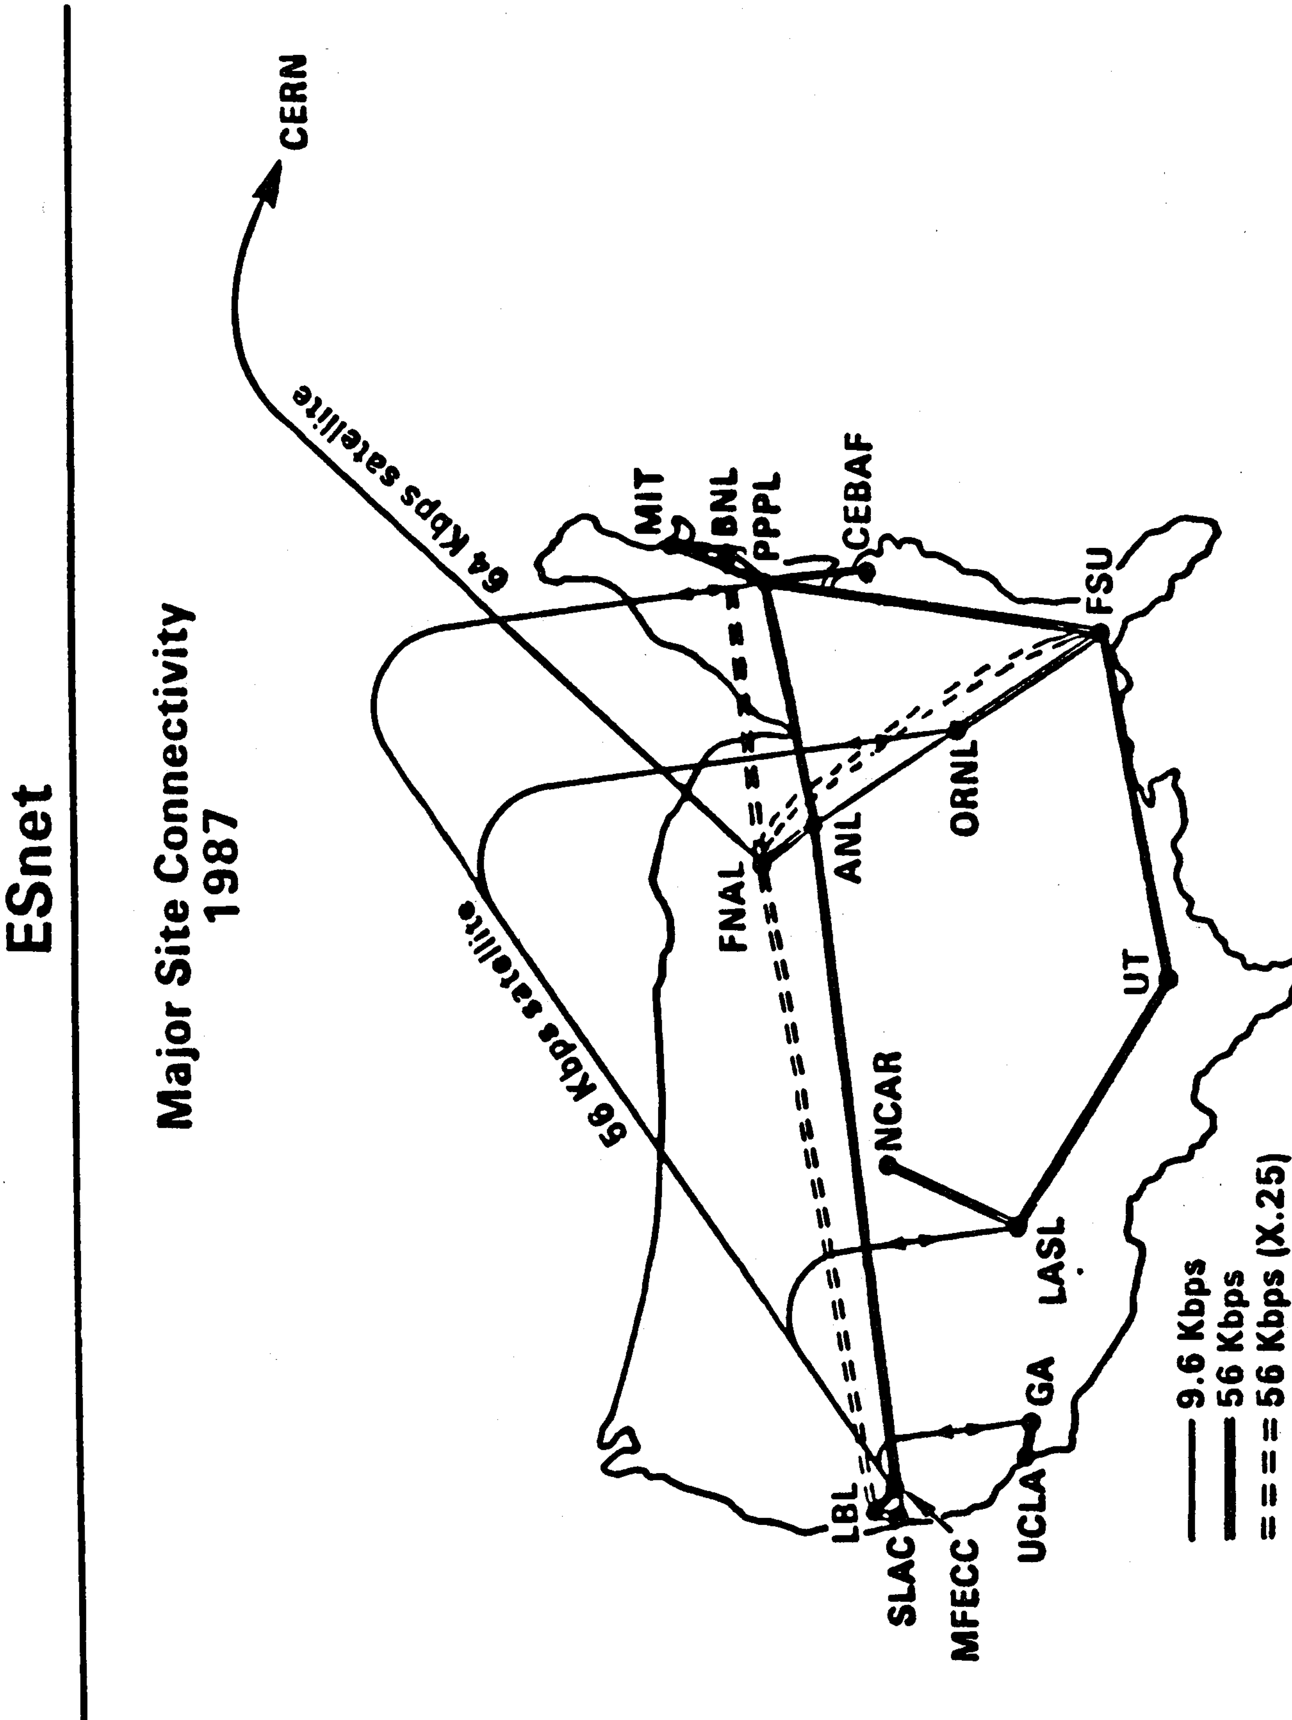
ESnet
Major Site Connectivity
1987
CERN
64 Kbps satellite
56 Kbps satellite
MIT
BNL
PPPL
CEBAF
FSU
ORNL
ANL
FNAL
UT
NCAR
LASL
GA
UCLA
LBL
SLAC
MFECC
9.6 Kbps
56 Kbps
56 Kbps (X.25)

Figure VI-2

CONNECTIVITY REQUIREMENT FOR ER COMPUTER NETWORK

GEOGRAPHIC LOCATIONS

This list of geographic locations has been compiled from ER program individual requirements documents. The locations are footnoted to designate 1.) that connectivity to this location is required by more than one ER program, 2.) that there are already multiple ER funded computer network links to the location, and 3.) that the location is serviced by or is part of a computer network funded by another agency, such as NSF, DARPA, etc.

ABILENE CHRISTIAN
ALABAMA 3
AMES LABORATORY 1, 2, 3
ARGONNE NATIONAL LABORATORY 1, 2, 3
ARIZONA 1, 3
ARIZONA STATE 3
ARKANSAS 3
ATLANTA
AUBURN
BALL STATE
BATTELLE PACIFIC NORTHWEST LABORATORY
BOSTON 1, 3
BPA-PORTLAND
BRANDEIS
BRIGHAM YOUNG UNIVERSITY 1
BROOKHAVEN NATIONAL LABORATORY 1, 2, 3
BROOKLYN COLLEGE
BROWN 3
CALIFORNIA/BERKELEY 2, 3
CALIFORNIA/DAVIS 1, 3
CALIFORNIA INST OF TECH 1, 2, 3
CALIFORNIA/LOS ANGELES 1, 2, 3
CALIFORNIA POLYTECHNIC STATE
CALIFORNIA/RIVERSIDE 1
CALIFORNIA/SAN DIEGO 3
CALIFORNIA/SANTA BARBARA 1, 3
CALIFORNIA/SANTA CRUZ 3
CARNEGIE-MELLON 1, 3
CEBAF
CHICAGO 1
CINCINNATI 1
CLARK
CLARKSON
COLORADO 3
COLORADO STATE 3
COLUMBIA 1, 3
CONNECTICUT 3
CORNELL 1, 3

Figure VI-3

DENVER 3
DEPT OF ENERGY-GTN 1, 2
DEPT OF INTERIOR-DENVER
DREXEL
DUKE 1, 3
ELMHURST
EMORY
ENVIRONMENTAL MEASUREMENT LAB-NY
FERMI NATIONAL ACCELERATOR LABORATORY 1, 2, 3
FLORIDA 3
FLORIDA STATE 1, 2, 3
FLOW RESEARCH
FRANKLIN/MARSHAL COLLEGE
GA TECHNOLOGIES 1
GEORGE MASON
GEORGE WASHINGTON 3
GEORGIA 3
GEORGIA TECH 3
GRUMMAN AERO
HANFORD 1, 2
HARVARD 1, 2, 3
HOUSTON 1, 3
IDAHO
IDAHO NATIONAL ENGINEERING LABORATORY 3
ILLINOIS-CHAM 1, 3
ILLINOIS-CHI 1, 3
ILLINOIS STATE 3
INDIANA 1
INHALATION TOXICOLOGY RES INST-ALBUQUERQUE
INST MINING & TECH-SOCORRO/NM
IOWA 1, 3
IOWA STATE 1, 3
JOHNS-HOPKINS 1, 2, 3
KANSAS 1, 3
KANSAS STATE 1, 3
KENT STATE
KENTUCKY 1, 3
LAWRENCE BERKELEY LABORATORY 1, 2, 3
LAWRENCE LIVERMORE NATIONAL LABORATORY 1, 2, 3
LEHIGH 1, 3
LOS ALAMOS NATIONAL LABORATORY 1, 2, 3
LOUISIANA STATE 1, 3
LOUISVILLE 3
MARYLAND 1, 2, 3
MASSACHUSETTS INST OF TECH 1, 2, 3
MCDONNELL DOUGLAS
MICHIGAN 1, 3
MICHIGAN STATE 1, 3
MINNESOTA 1, 3
MISSISSIPPI 3
MISSOURI 3
MONTANA

NASA/LEWIS RESEARCH CENTER
NATIONAL BUREAU OF STANDARDS-JILA
NCAR 1, 3
NEBRASKA 3
NEW HAMPSHIRE 1
NEW MEXICO 1
NEW MEXICO IMT
NEW MEXICO STATE
NEW YORK 1, 2, 3
NEW YORK MEDICAL CENTER 1
NEW YORK, STATE UNIV/STONY BROOK 1, 2, 3
NORTH CAROLINA/CHAPEL HILL 1, 3
NORTH CAROLINA STATE 1, 3
NORTHEASTERN 3
NORTHRIDGE
NORTHWESTERN 1
NOTRE DAME
OAK RIDGE NATIONAL LABORATORY 1, 2, 3
OHIO 1
OHIO STATE 1, 3
OKLAHOMA 3
OREGON
PACIFIC NORTHWEST LABORATORY 1, 2
PENN 1, 3
PENN STATE 1, 3
PHILADELPHIA
PITTSBURGH 3
PRINCETON 1, 3
PRINCETON PLASMA PHYSICS LABORATORY 1, 2
PURDUE 1, 3
RENSSELAER POLYTECHNIC INSTITUTE
RICE 1, 3
ROCHESTER 3
ROCKWELL HANFORD
RUTGERS 1, 3
SANDIA NATIONAL LABORATORY 1, 2, 3
SAN FRANCISCO STATE
SCIENCE APPLICATIONS, INC 1, 2, 3
SCRIPPS INSTITUTION OF OCEANOGRAPHY
SKIDAWAY INST OF OCEANOGRAPHY
SOLAR ENERGY RESEARCH INST
SOUDAN
SOUTH CAROLINA
SOUTHEASTERN
SOUTHWEST RESEARCH INSTITUTE
ST. LOUIS
STANFORD 1, 3
STANFORD LINEAR ACCELERATOR CENTER 1, 2, 3
SYRACUSE 1, 3
TEMPLE
TENNESSEE 3
TEXAS 1, 2, 3
TEXAS A&M 1, 3

TEXAS ACCELERATOR CENTER 1, 3
TEXAS TECH 3
TOLEDO
TRW INC
U.S. AIR FORCE ACADEMY
U.S. NAVAL ACADEMY
UTAH
UTAH STATE
VALPARAISO
VANDERBILT 3
VASSAR
VIRGINIA 3
VIRGINIA POLY INST 3
VIRGINIA STATE
WASHINGTON 3
WASHINGTON STATE
WILLIAM & MARY
WISCONSIN 3
WORCESTER POLYTECHNIC INSTITUTE
WYOMING
YALE 3

VII. Foreign Access Requirements

The Energy Sciences Steering Committee has recognized the need for European and Japanese data communication links to support ER program involvement in several recognized international research collaborations. ER is currently experiencing a phenomenal growth in these international requirements, as large scientific collaborations have become increasingly international in scope. At present, ER scientists collaborate on experiments and projects at research facilities in Switzerland, Canada, Italy, England, Japan, and Germany. Of the existing collaborations, most of the requirements have been in the Fusion and High Energy and Nuclear Physics programs. However, new requirements in Basic Energy Sciences and Health and Environmental Research are forming along with significant requirements increases in the existing collaborations.

The Scientific Computing Staff (SCS) has been tasked with the responsibility of implementing a nationwide, multi-program network for research. To accomplish this, the SCS formed the ESNET steering committee which has recognized these foreign requirements as a significant factor in meeting the goals of ESNET. It is the consensus of the ESNET steering committee to implement foreign data links from major nodes on the ESNET backbone in order to provide easy access by all ER researchers to other collaborators and to the facilities themselves. [8] The steering committee has agreed to pursue consolidated links to all foreign locations that serve multi-program requirements to promote cost effectiveness in order to provide broader access to foreign networks and facilities.

Currently ER programs have identified requirements to collaborate with the following locations:

GERMANY

Deutshes Elektron-Synchrotron
Hamburg

Max-Planck-Institute
Munchen

University of Wuppertal
Wuppertal

Albert-Ludwigs-University
Freiburg

European Molecular Biology Laboratory
Heidelberg

KFA Juelich
Juelich, FRG

University of Stuttgart
Stuttgart, FRG

KFA Karlsruhe
Karlsruhe, FRG

Berlin Electron Synchrotron Light Source
Berlin

Gesellschaft fur Schwerionenforschung, GSI
D-6100, Darmstadt, West Germany

ENGLAND

JET, Culham Laboratory
Culham/Harwell
Daresbury (UK)

Rutherford Appleton Laboratory
Chilton, Didcot, Oxon

FRANCE

Cadarache - CEN

ITALY

Frascati - Frascati Energy Research Center - ENEA
Padova - IGICGR - ENEA
Gran Sasso Lab

AUSTRIA

IAEA Headquarters
Vienna

SWITZERLAND

CRPP, EPFL - Luasanne
CERN, Geneva
SIN, Villigen

JAPAN

Kyoto Univ, Uji - Kyoto
IPP - Nagoya
JAERI - Naka
JAERI - Tokai
KEK National Laboratory for High Energy Physics - Tsukuba
National Institute of Genetics - Misima, Sizuoka-Ken

Foreign data links are being operated in Fiscal Year 1987 by the Fusion Energy (FE) and High Energy Physics Program (HEP). HEP is currently operating low speed links between MIT and CERN, Fermi National Laboratory and Italy, and Lawrence Berkeley Laboratory and Japan. FE currently is operating data links between Oak Ridge National Laboratory (ORNL) and Germany [8], ORNL and England, and Lawrence Livermore National Laboratory and Japan. It should be noted that the data link installed during Fiscal Year 1987 between Italy and Fermi National Laboratory is being paid for by the Italians. The more common practice, with the exception of the CERN Laboratory in Switzerland, is for foreign collaborators to share costs for international service by paying for their "half circuit" in their respective countries. The attached reference documents address the Fiscal Year 1987 program requirements.

Plan of Action

The ESNET steering committee recommended the following actions to satisfy the ER foreign access requirements based on existing programmatic needs.

1. Proceed with a 64 Kbps satellite link from Fermi National Laboratory to CERN Laboratory in Switzerland in conjunction with planned interim upgrades to 56 Kbps in U.S.

2. Proceed with an interim 9.6 Kbps link from Lawrence Berkeley Laboratory to KEK in Tsukuba, Japan while continuing to pursue a more cost effective, longer term solution, i.e., a consolidated high speed data link to Japan that would service all Energy Research programs in the U.S.

3. Initiate planning for consolidation of all German requirements through a single high speed wide band satellite data link.

4. Review and evaluate requirements with Italy.

5. Assess and evaluate new requirements as they occur.

6. Proceed, during FY 1987, with a network security study of the Fermi National Accelerator Laboratory data link to CERN and expand this study to other data link requirements in Fiscal Year 1988.

VIII. Implementation Plans

There are many permutations for implementing or expanding a computer network such as ESNET. These mainly deal with varying the degree of deployment of various components such as hardware and communication lines, proposing alternative deployment or upgrade schedules, and selectively implementing various functional requirements. Alternatives in this regard will be considered in the narrative discussion of each individual network component and will not be presented as totally different and complete scenarios.

The recommended budget and implementation scenario, which is presented below, assumes a modest implementation schedule. It meets interagency commitments for interoperability in the near term, it also meets the deadline for recompetition of the existing MFENET communication media contract, and it recognizes the ER user needs as expressed in the preceding Requirements Section.

The consensus of the ESNET Steering Committee is that the ESNET facility will be a valuable asset for assisting research which will provide uniform benefits and will be available to all ER programs. Therefore, the ESNET Steering Committee agreed that the ESNET project should be endorsed by all ER programs and that ER should vigorously seek additional funding to support the ESNET project. It is important to note that the total recommended budget presented here for the ESNET initiative amounts to less than 1% of the total ER annual budget. This is very small relative to the potential payback for such a facility. It is commensurate with those reported for data communications expenditures in industry publications and significantly less than the NSF request for $46,000,000 for computer network support for FY 1989.

Figure VIII-1 presents the proposed three year incremental budget for the ESNET initiative. This budget is termed incremental because it is that amount needed, over and above existing funds for ER supercomputer access support, to expand computer network coverage across ER programs. There are also funds currently being expended within research grants for data communication support (e.g., the 9600 baud line from Brandeis to Harvard in the HEP program under $10,000/year) which are individually low cost, but which may form a modest aggregate sum. As these facilities are upgraded or incorporated into the ESNET, this funding should remain in the program areas for continued research support for a couple of reasons. For the most part, these funds have displaced other necessary research support due to the importance of providing some connectivity, although minimal, for collaboration support. Also, these existing facilities are essential and must be phased in through a well managed process which spans fiscal year boundaries and which is not conducive to reprogramming or other fiscal procedures.

The first network infrastructure component of the proposed ESNET initiative is personnel. The central facility at the NMFECC will need to increase the network staff as the community which the NMFECC services grows and as the demand for more functionality and control room operations increases. This budget supports two additional FTE in FY 1988 and another three additional

FTE by the end of FY 1990. Another vital part of the network infrastructure is capable support at major ESNET nodes as discussed in ESNET Infrastructure section. It is assumed that the requirement at each site will vary with the complexity of local area network connections, with the number of other outside connections, with the number of different hardware interfaces, etc. Hence, some but not all major sites will require a full time FTE. The proposed budget assumes growth to six full and six half time FTE in 1989.

ESNET 3 YEAR INCREMENTAL PROPOSED BUDGET
($X1000)

NETWORK INFRASTRUCTURE COMPONENT	FY 1987	FY 1988	FY 1989	FY 1990
Personnel				
At NMFECC	0	200	325	450
At Major ESNET Nodes	0	725	950	950
Hardware				
ESNET routing - NMFECC supplied (LTO)	0	0	300	600
- Maintenance	0	0	75	150
Gateways to other networks - install	0	125	250	250
- operate	0	25	100	150
X.25 migration - 5 sites inc. capital, maint., and misc	100	100	100	75
Software	0	200	250	300
Communication Lines				
Expanded Backbone	50	500	500	500
Multi-megabit 2 way broadcast	0	0	500	1200
Vsat 1-way multicast	0	0	400	1000
Expand/upgrade tail circuits	0	200	500	1000
Terrestrial T-1 prototype	0	0	410	800
Foreign data links	0	600	1200	1200
Research and Development	0	300	600	900
Support and Other Services				
Workshops	0	100	100	100
Newsletters, travel, etc.	0	175	225	300
TOTALS	150	3250	6785	9925

Figure VIII-1

The second component is hardware, including NMFECC supplied ESNET IP hardware, gateways to other agency networks, and hardware needed to support X.25 in the near term. The main budgetary consideration in this area is the extent of deployment. The assumption is that ESNET will acquire the ESNET IP hardware via a lease-to-ownership (LTO) with incremental delivery over several years. In this fashion, fiscal outlays can be fixed while year to year and total quantities can vary. The total needed is between 50 and 100 hardware nodes to be provided by an annual LTO outlay of $600,000 for either 30 or 60 months, respectively.

In light of geographic coverage by other networks, such as NSFNET (over 70% duplication), it would be very cost advantageous to ER to utilize these other networks to provide network access to ER supported researchers at most of these duplicative locations. However, flexibility is necessary to provide additional ESNET hardware at those locations where expected or desired performance levels are not achieved. The gateways to other networks included in this proposed budget would provide this interoperability with these other networks. Current plans are to install two such gateways per year for the planning period.

The last hardware budget item is the support for X.25 computer network protocols for two years as part of the ESNET migration plan for supporting the large DECNET requirement of the ER research community. The planned and budgeted ESNET hardware will provide flexibility and will handle X.25 protocol interfaces as well as IP and others. However, the support for X.25 is needed now and must be supported in the interim while the ESNET software is being implemented. This X.25 migration plan has received consensus support by the ESNET steering committee and has already been initiated during FY 1987.

The next network component is software. This inclusion is for commercial software products and is important in many areas, especially in products such as vendor supplied TCP/IP products and application level packages for implementation throughout the community.

Communication lines, the next network component, are the vital physical links between sites which carry data/information. The breakdown here includes incremental costs to add additional sites to the current MFENet backbone for the ESNET, especially for the X.25 migration plan. The proposed budget also includes adding service to approximately 25-30 locations for high speed 2-way broadcast and for the multicast technology during the planning period. This number of locations is modest relative to the total number of locations identified in the Requirements Section because of the anticipated utilization of providing connectivity to many of these locations via gateways to other networks. Nonetheless, approximately 40-50 locations (network tail circuits) will not have other network connectivity or will require upgrades to meet performance objectives. Hence, the additional deployment or upgrade for these terrestrial tail circuits is included as the "expand/upgrade tail circuits" item to accommodate growth to this number of locations by the end of FY 1990. The terrestrial T-1 prototype is included to support four backbone communications links for research and development of type of service routing, especially involving distributed graphics, visualization, and other high performance interactive projects. The foreign data links are included as per the requirement detailed in the Foreign Access Requirements Section. This includes six links, assuming that the link between FNAL and CERN is continued to be funded through FSU and that the Japan and German data links are consolidated into the ESNET process.

Research and development is an essential component of ESNET. It ensures that future Energy Research will continue to employ the most advanced computational networking facilities in the world. Network research and development has two basic goals: increased functionality and increased performance. Increased functionality creates opportunities for scientific breakthroughs by bringing widely dispersed supercomputing and data resources into the user's local workstations from data generated by a remote supercomputer, or users may edit remote data files using a highly interactive editor that runs locally. Such high technology facilities create a test bed for advanced networking products. This gives U.S. industry a competitive edge over foreign suppliers. A recent interagency computer network study found that to support this high level of networking technology, a thousand fold increase in network capacity is needed. Providing this increase at an affordable cost is what motivates research in increased performance. Areas to be investigated include congestion detection and control algorithms, and specialized data compression hardware and software. These techniques increase the effective capacity of existing communications lines, thus increasing capacity without increasing cost. Other areas of investigation include high speed, multi-media services, and user information services. Three or four small projects are planned for initiation in FY 1989.

Interagency coordination ensures compatibility between ESNET and other government agency networks such as the ARPA/Milnet and NSFnet. It is critical to ER programs that such compatibility exist and be maintained. The resultant cost savings have been detailed earlier in this document. The establishment of network and product standards is equally critical to the success of ESNET within the ER community, especially where interagency or international collaborations are necessary. The line budget for this item reflects staffing and travel costs for efforts expended for standardization meetings and interagency discussions.

The last computer network infrastructure component which is included is for information dissemination and assistance services which will help bind the ESNET user community together, which will promote fast problem resolution, which will provide valuable user input for network direction, and which will promote synergism throughout the entire ER computer network community.

To summarize the progression or status of the proposed ESNET implementation at the end of each fiscal year, as presented in Figure VIII-1, will be as follows:

FY 1987

The ESNET implementation will begin during FY 1987 with the completion of the ESNET design and software development (no deployment at this point), with the initiation of the X.25 migration plan (installation of extended backbone connections and X.25 data switches as described in reference [9]) and with the endorsement of the ESNET program plan for inclusion in the DOE budget.

FY 1988

At the end of FY 1988, the ESNET infrastructure should be well established. Personnel will be added to the network software staff at the NMFECC and also to the staffs at about nine, assuming five full time and four half time staff, of the ESNET major nodes. The X.25 migration plan will be fully implemented and X.25 support as part of the ESNET routing hardware and software will begin to be tested. Gateways to other networks will be under development for critical connections such as to the group of networks comprising the NSFNET. It is anticipated that by the end of FY 1988, at least two gateways will be installed to these networks, such as to the NSF backbone and to the SURANET. With regard to foreign data links, three consolidated connections will be initiated for large requirements such as CERN, West Germany, and Japan. Proposed funding for FY 1988 will allow the initiation of three small network research projects. Lastly, during FY 1988, two workshops for ER computer networks and applications will be conducted and ESNET introductory documentation and newsletters will be distributed throughout the research community.

FY 1989

By the end of FY 1989, personnel will be added to supplement the network operations staff at the NMFECC and also to provide ESNET support at three more major nodes (two full time and one half time FTEs). ESNET routing hardware will be deployed throughout FY 1989 to approximately 20-25 non-supercomputer supported locations, gateways will be implemented for two more outside networks such as ARPANET and another NSF network (total gateways will be at 4), and X.25 migration support will continue as ESNET supported X.25 routing begins deployment. With regard to communications lines during FY 1989, the proposed ESNET budget will continue support for the expanded ESNET backbone at 56 Kbps, will allow for approximately 30 upgrades of tail circuits to 56 Kbps and will allow for consolidated support for three more foreign data links, such as Italy, England, and Canada. This budget will also include funding to initiate development and deployment for Very Small Aperature Satellite (VSAT) technology to support high speed broadcast and multicast capabilities at about 10 major ESNET nodes and to prototype terrestrial T-1 circuits (preferably fiber optic circuits) through a cross country segment of the ESNET backbone, such as LLNL-ANL, ANL-PPL, and PPL-MIT. Several more research projects will be initiated and workshop, training and documentation support will continue.

FY 1990

During FY 1990, personnel support will level off with the addition of one FTE for network operations at the NMFECC. Also, funding for ESNET routing hardware and gateway support will level off providing approximately 30-35 ESNET installations and two more gateways. The funding for the X.25 migration plan will discontinue, provided that the ESNET X.25 support is fully implemented during FY 1990. With regard to communication lines, it is anticipated that all costs in this area stabilize except for those for terrestrial T-1 circuits. These T-1 costs may increase in future years if these circuits prove fruitful for supporting ER network applications. Funding for computer network research projects will also level off at about ten percent of the ESNET budget. Finally during FY 1990, funding will continue for workshops and for ER computer network workshops and for ESNET training and documentation.

Reference 1

NETWORK REQUIREMENTS FOR SCIENTIFIC RESEARCH, 1987, LEINER

ATTACHED REFERENCES

1. Network Requirements for Scientific Research, 1987, Leiner

2. MFENET II Overview, 1986, Leighton

3. Energy Sciences Network Proposed Implementation Plan, 1986

4. Report on Interagency Networking for Research Programs, 1986, FCCSET

5. The Effect of Distributed Computing Technology on Wide Area Network Capacity Requirements, 1987, Hall, Johnston, Rosenblum

6. Functional Characteristics of ESNET, 1987, Dunning, Greenwald, Loken

7. HENP, FE, AMS, HER, and BES Program Requirements, 1987

8. European Data-Link Assessment, 1987

9. ESNET X.25 Migration Plan, 1987

Network Requirements for Scientific Research

Internet Task Force on Scientific Computing
Barry M. Leiner, ed.

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 87.1
January 1987

Computer networks are critical to scientific research. The recognition of that fact has prompted several agencies to fund networks for their researchers. This workshop is aimed at investigating the cooperation between these agencies to in order to provide these functions to the broad scientific community in a cost-effective manner. This paper attempts to outline the requirements for such a national research internetwork. It first addresses the functions a user requires of a network. It then addresses near term requirements and future goals for such a network.

This is a task force of the Internet Activities Board, and is chartered to investigate advanced networking requirements that result from scientific applications.
Work reported herein was supported in part by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Adminstration (NASA) to the Universities Space Research Association (USRA).

Network Requirements for Scientific Research

Internet Task Force on Scientific Computing
Barry M. Leiner, ed.

Research Insitute for Advanced Computer Science
NASA Ames Research Center

January 1987

1. INTRODUCTION

Computer networks are critical to scientific research. They are currently being used by portions of the scientific community to support access to remote resources (such as supercomputers and data at collaborator's sites) and collaborative work through such facilities as electronic mail and shared databases. There is considerable movement in the direction of providing these capabilities to the broad scientific community in a unified manner, as evidenced by this workshop. In the future, these capabilities will even be required in space, as the Space Station becomes a reality as a scientific research resource.

The purpose of this paper is to identify the range of requirements for networks that are to support scientific research. These requirements include the basic connectivity provided by the links and switches of the network through the basic network functions to the user services that need to be provided to allow effective use of the interconnected network. The paper has four sections. The first section discusses the functions a user requires of a network. The second section discusses the requirements for the underlying link and node infrastructure while the third proposes a set of specifications to achieve the functions on an end-to-end basis. The fourth section discusses a number of network-oriented user services that are needed in addition to the network itself. In each section, the discussion is broken into two categories. The first addresses near term requirements: those capabilities and functions that are needed today and for which technology is available to perform the function. The second category are long term goals: those capabilities for which additional research is needed.

2. NETWORK FUNCTIONS

This section addresses the functions and capabilities that networks and particularly inter-networks should be expected to support in the near term future.

2.1. Near Term Requirements

There are many functions that are currently available to subsets of the user community. These functions should be made available to the broad scientific community.

2.1.1. User/Resource Connectivity

Undoubtedly the first order of business in networking is to provide interconnectivity of users and the resources they need. The goal in the near term for internetworking should be to extend the connectivity as widely as possible, i.e. to provide ubiquitous connectivity among users and between users and resources. Note that the existence of a network path between sites does not necessarily imply interoperability between communities and or resources using non-compatible protocol suites. However, a minimal set of functions should be provided across the entire user community, independent of the protocol suite being used. These typically include electronic mail at a minimum, file transfer and remote login capabilities must also be provided.

2.1.2. Home Usage

One condition that could enhance current scientific computing would be to extend to the home the same level of network support that the scientist has available in his office environment. As network access becomes increasingly widespread, the extension to the home will allow the user to continue his computing at home without dramatic changes in his work habits, based on limited access.

2.1.3. Charging

The scientific user should not have to worry about the costs of data communications any more than he worries about voice communications (his office telephone), so that data communications becomes an integral and low-cost part of our national infrastructure. This implies that charges for network services must NOT be volume sensitive and must NOT be charged back to the individual. Either of these conditions forces the user to consider network resources as scarce and therefore requiring his individual attention to conserve them. Such attention to extraneous details not only detracts from the research, but fundamentally impacts the use and benefit that networking is intended to supply. This does not require that networking usage is free. It should be either be low enough cost that the individual does not have to be accountable for "normal" usage or managed in such a manner that the individual does not have to be concerned with it on a daily basis.

2.1.4. Applications

Most applications, in the near term, which must be supported in an internetwork environment are essentially extensions of current ones. Particularly:

Electronic Mail

Electronic mail will increase in value as the extended interconnectivity provided by internetworking provides a much greater reachability of users.

Multimedia Mail

An enhancement to text based mail which includes capabilities such as figures, diagrams, graphs, and digitized voice.

Multimedia Conferencing

Network conferencing is communication among multiple people simultaneously. Conferencing may or may not be done in "real time", that is all participants may not be required to be on-line at the same time. The multimedia supported may include text, voice, video, graphics, and possibly other capabilities.

File Transfer

The ability to transfer data files.

Bulk Transfer

The ability to stream large quantities of data.

Interactive Remote Login

The ability to perform remote terminal connections to hosts.

Remote Job Entry

The ability to submit batch jobs for processing to remote hosts and receive output.

Applications which need support in the near term but are NOT extensions of currently supported applications include:

Remote Instrument Control

This normally presumes to have a "human in the control loop." This condition relaxes the requirements on the (inter)network somewhat as to response times and reliability. Timing would be presumed to be commensurate with human reactions and reliability would not be as stringent as that required for completely automatic control.

Remote Data Acquisition

This supports the collection of experimental data where the experiment is remotely located from the collection center. This requirement can only be satisfied when the bandwidth, reliability, and predictability of network response are sufficient. This cannot be supported in the general sense because of the enormous bandwidth, very high reliability, and/or guaranteed short response time required for many experiments.

Reference 2

MFENET II OVERVIEW, 1986, LEIGHTON

How It Will Be Implemented

Figure 1 illustrates the planned architecture of MFEnet II. The network will comprise communication control processors (CCPs), gateway processors, and server processors. CCPs manage the long-haul communication facilities such as satellite links, and are responsible for data routing, diagnostics and statistics, and downloading processor software. The gateway processors provide a gateway to the network for the user. The gateways will provide access compatible with the Internet Protocol (IP) used by the DARPA Internet. Additional user access levels may be supported later, including X.25 and/or ETHERNET. The server processors provide various functions for the user, including terminal access, printer/plotter output, and some specialized graphics support.

The computer hardware will be updated through a competitive procurement to select a vendor of a 32-bit computer engine and associated peripheral equipment. A larger, but compatible, development system will also be procured for network software development.

NMFECC's CRAY computers already use the NSP protocol extensively and will continue to do so. The IP protocol will be added to the CRAYs to allow connection to the network gateways. The server processors will also use NSP, and VAX software will be provided to allow NSP access to the CRAYs.

We are currently developing plans to implement the TCP protocol on the CRAY computers to support TCP/IP access. Our plans are to support both NSP and TCP concurrently until the move to standard protocols can be completed.

Capabilities of the New Network

The new network architecture will support these additional capabilities:

- **ISO Standard Protocols:** This new network structure is a move toward easy substitution of ISO standard protocols as they become available. The IP protocol at the gateway will be replaced with (or additionally support) the ISO equivalent, also known as IP. Since the network is not sensitive to higher level protocols, substitution of the ISO transport layer "TP4" and higher levels can be accommodated without change in the network.
- **Inter-networking:** Because the network will be compatible with the DARPA internet protocol, inter-networking will be easier. These two networks will be able to connect gateway-to-gateway, and serve as "data carriers" for each other. Additionally, these gateways provide a well defined demarcation point for accounting and access control, if required, for such interagency activities. A draft report from a recent meeting reviewing internetworking issues stated that this plan provides "... a significant step toward functional interoperability between the most successful networking technologies serving the scientific research community."

The network design will allow use of the commercially available TCP/IP protocol set. This will allow multi-vendor use of the network for services supported by this protocol structure. This implies that the multitude of workstations and computers that use TCP/IP will be able to use MFEnet II for data communications.

The additional effort of adding TCP to NMFECC's CRAYs will allow the TCP/IP community, as well as the NSP-based MFE and ER communities, to use the CRAYS remotely.

Additional work is necessary to support the DECnet-based communities, but this network architecture allows several approaches, such as X.25 for network access. This level of support will allow DECnet users to use MFEnet II as a "transparent" data communications carrier for DECnet interactions.

We have no plans to provide direct DECnet support on NMFECC 's CRAYs. Additional work, however, is currently under way to provide NSP service capabilities to DECnet users. This will provide NMFECC CRAY access from a DECnet host with MFEnet/NSP user level interaction.

- **Common Data Communications Facility:** The new network structure clearly delineates the network interface from the host and server access software. Inclusion of the IP compatible gateway allows use of the network as a common data communications facility. The addition of X.25 as an access level will enhance and further generalise this facility.
- **LAN Access for ETHERNET and IP-based Computers:** ETHERNET has been chosen as the "link level" interface for the gateway. This in conjunction with the ETHERNET-based servers will provide access via the user's LAN. The same ETHERNET is, of course, capable of providing access for IP-based host computers. Host computers using either TCP/IP or NSP/IP will also have access to the network services.
- **International Communications:** Many European sites interested in international communications with the United States ER community are DECnet users. A multitude of other protocols, however, are in use. These include the JANET "coloured books," the IBM Remote Spooling Communication System (RSCS), TCP/IP, and versions of "home grown" protocols. Data communications can probably best be facilitated in the short term by the addition of X.25 to MFEnet II. In the longer term, both the European community and we are committed to supporting ISO standard protocols. This may be coordinated so that these efforts can be brought together as soon as feasible.

Transition

Our plan is to physically construct the MFEnet II in parallel with the current network by incrementally phasing individual sites from one network to the other. Since the conversion is expected to take three years, users will be split during this period between the two networks (although both will provide access to NMFECC). This problem will be handled by careful planning and by simultaneously moving user sets with common communication requirements. One or more gateways will also be required that will support communication between the networks in the interim. These gateways may use manual forwarding techniques, or may require automatic forwarding.

Schedule

We have begun a competitive procurement process to acquire the 32-bit CPU engines required for this new network. The hardware procurement and deployment schedule will be as follows:

1 May 1986	RFQ released
15 July 1986	RFQ responses due
1 August 1986	RFQ responses evaluated, vendor selected
1 November 1986	Delivery of Initial Configuration of Equipment
Second Quarter FY87	Completion of first installation
Fourth Quarter FY89	Network change-over completed

The software development is under way now. The initial capability will be available in the second quarter of FY87.

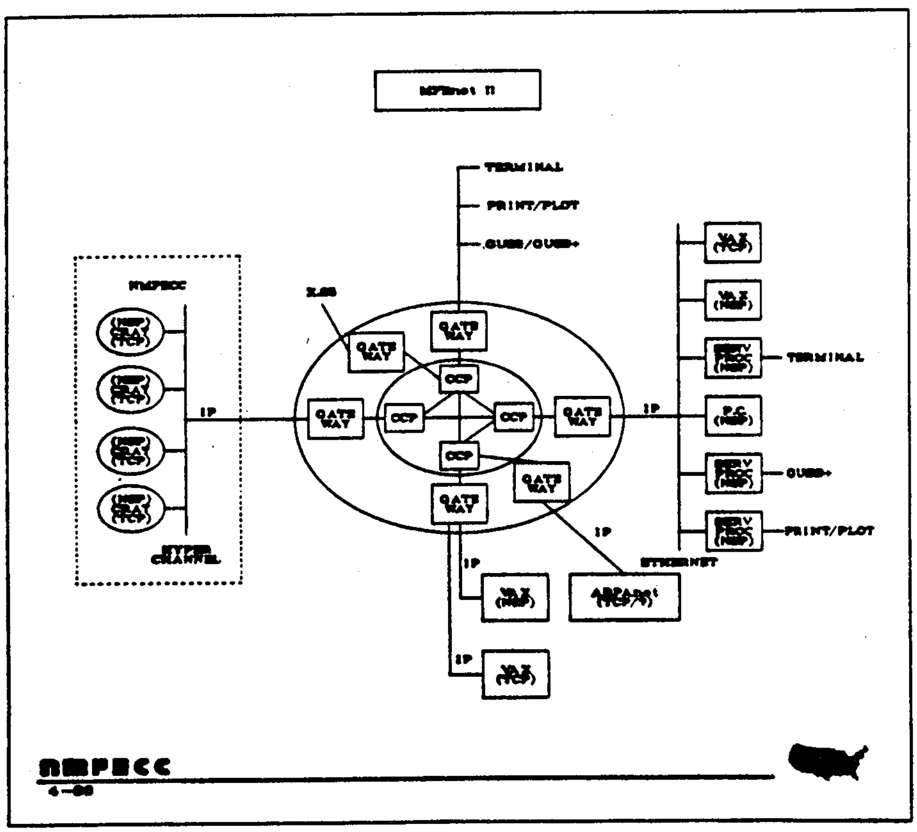

Figure 1. Proposed architecture of MFEnet II.

Reference 3

ENERGY SCIENCES NETWORK PROPOSED IMPLEMENTATION PLAN, 1986

ENERGY SCIENCES NETWORK

PROPOSED IMPLEMENTATION PLAN

INTRODUCTION

During FY 1985, the Scientific Computing Staff was directed to evaluate the status of, and the requirement for, Energy Research (ER) data communications as an adjunct to its responsibility for providing remote access to ER's supercomputer facilities. The various ER programs have begun to recognize the importance of data communications in support of remote computing, nationwide and international collaborations, software research and development for remote projects and data file transfers. For example, HER plans to initiate a centralized software activity which will be remotely accessed via a common ER data communications facility. A general purpose ER Network concept has begun to emerge with a proposal to redesign the existing MFENET as described in a recent paper by Jim Leighton (NMFECC Buffer, May 1986). This new concept has been reviewed by the interagency internet community (DARPA, NSF, NASA, etc.) and was determined to be a major step toward the creation of an interagency research internet.

MANAGEMENT

The migration to an Energy Sciences Network (ESNET) should be evolutionary in nature so that current, critical requirements are addressed and so that overall benefits of operating a single ER network can be achieved. The proposed approach will be to combine the current network activities of the various ER programs by coordinating the applications level requirements with the SCS and by managing the network level activities through the SCS. These network level activities will be advised through a steering committee composed of representatives of ER program with the help of technical consultants as required. Installation, coordination and operation of the data communications facilities are proposed to be the responsibility of the MFECC network staff at LLNL. Centralized management of the proposed ESNET by the SCS should provide for effective control and operation, while the proposea steering committee will assure that long term goals are achieved without impacting ongoing requirements.

OPERATIONS

It is proposed that the networking staff of the MFECC assume the responsibility for coordinating the requirements as determined by the steering committee, for planning and installing new data communications facilities as required, and for the daily operation of the resulting ESNET.

This additional network level activity will be handled initially by an additional 2-3 staff members at the MFECC.

IMPLEMENTATION

1st Phase - The initial activity would be concerned with implementing the above recommendations and getting the procedures, committee, and support staff in place. The additional and/or special requirements of the HEP, BER, and other ER groups will be reviewed during this period, and appropriate responses to their needs will be determined.

Although the long range goal is to combine data communications activities in a cost effective manner, for this first phase, this will be practical in only a limited manner.

This phase will probably require 12-18 months.

2nd Phase - This phase will be the beginning of integration of data communications requirements onto common facilities.

The MFEnet II proposal will have progressed sufficiently to address the additional requirements of new user communities, at this point all ER requirements will begin to be merged and use of the MFEnet will be considered where practical.

Reference 4

REPORT ON INTERAGENCY NETWORKING FOR RESEARCH PROGRAMS, 1986, FCCSET

REPORT

ON

INTERAGENCY NETWORKING FOR RESEARCH PROGRAMS

BY THE

FCCSET COMMITTEE ON VERY HIGH PERFORMANCE COMPUTING

February 19, 1986

Introduction

In February 1985, the FCCSET Committee on Very High Performance Computing (VHPC) established a Networks Working Group to exchange information on the various Agency plans for the development of Supercomputer Access networks.

The Networks Working Group reported in October 1985. The report recommended that an interagency task force be established to examine the requirement and justification for a well coordinated, encompassing, computer networking infrastructure for supercomputer access and research collaboration in the United States, and that this interagency task force should develop a plan for implementing such a networking infrastructure. The primary issues highlighted in the report were those of avoiding redundancy, maximizing productivity, and enhancing collaborative research efforts by linking the existing networks together in a highly functional manner.

The Networks Working Group Report also suggested that the approach to be taken to building this networking infrastructure is to interconnect the existing networks into a network of networks through the use of highly functional network gateways. This approach was seen as the most economical method for providing inter-communication between research networks because it takes advantage of existing network facilities. In addition, the approach was seen as necessary to allow the individual agency networks to continue to utilize various networking standards, to meet their distinct requirements and operational considerations, and to allow cost recovery and access control for network use to remain a funding agency responsibility as directed through existing procedures.

The FCCSET VHPC Committee requested that the Networks Working Group develop the outline plan for the recommended networking infrastructure. In this Final Report, the Networks Working Group recommends the formation of an Interagency Research Internet organization to build the networking infrastructure, and outlines its functions, responsibilities and funding requirements.

The Interagency Internet Model

The interconnection of computer communication networks, so that users on a computer on any network can fully inter-communicate with users or computers on any other network as if they were working on their own network, creates a resultant "network of networks" which is generally referred to as an "Internet," in reference to the ARPA Internet model. This model provides the structure to allow the interconnection of networks, without demanding that each individual network be owned and controlled by the Internet organization. Each network in the Internet can be controlled and developed by a separate organization, with the interconnections between the network and any other network, 'regulated,' and therefore indirectly controlled and managed, by the Internet authority. To connect to the Internet, a set of gateway protocols and standards must be adhered to, and these protocols and standards are determined, in cooperation with the constituent network authorities, by the Internet management.

The Internet model is thus the appropriate model for managing, through active coordination, the proposed Interagency research network infrastructure.

The Interagency Research Internet Organization

The research network infrastructure will be coordinated by an "Interagency Research Internet" (IRI) organization. The IRI organization, under the direction of the FCCSET Committee on Very High Performance Computing (VHPC), will provide overall coordination of the management and operation of the Interagency Internet. The IRI will establish, promulgate, and coordinate through interagency cooperation, protocol standards and functional standards for interagency internetting. IRI will also address issues of documentation and information availability between the involved agencies. IRI will actively coordinate interagency internetting research projects being undertaken within the individual agencies, through proposal and status reviews. However, the resources required for internetting research projects will continue to be budgeted through the individual agency programs. The agencies participating will include: NSF, NASA, DOE, and DARPA. Other agencies may join at a later date. The IRI structure is modeled on the existing DARPA Internet structure.

Oversight:

Overall policy and goals for internetting will be developed by IRI under the direction of the FCCSET VHPC Committee. IRI will be comprised of members from those Federal Agencies which currently operate networks in support of research activities. The chairperson of IRI will be appointed by the FCCSET Committee. Day to day management of internetting activities will be the responsibility of the network project managers within the individual agencies. These activities will be actively coordinated by IRI through periodic proposal and status reviews and through the annual establishment of internetting plans and goals. IRI will also be responsible for protocol and functional internetwork standards.

Management and Operation:

Activities to create and support an interagency research network infrastructure include the purchase, installation, and operation of the major or 'core' network gateways between the existing and planned Research Networks. It will certainly involve the software development and maintenance, hardware maintenance, and operational monitoring and control of these gateways so that the Interagency Internet is an available and reliable communications entity. Activities will also include the installation of network routing, access control, and accounting procedures and tools, as these are developed.

Internetwork Development:

This area will involve identifying the research and development projects necessary to create, maintain, and enhance the Inter-agency Internet. IRI will coordinate the establishment and implementation of these projects with the constituent research networks. The actual projects to be undertaken will be agreed to by IRI and the participating agencies, and will be implemented by the constituent agencies or their contractors.

Approach:

A phased approach should be undertaken toward the creation of an Interagency Research Internet. The interconnectivity will be based on the use of common protocols and on the use of IP Gateway technology. The initial phase will place emphasis on connectivity rather than on performance between networks. It will be marked by the consummation of Memoranda of Understanding between agencies and by a movement toward common protocols and applications gateways. The second phase will emphasize experi-mentation and research in paradigms and in the management and control of distributed networked systems. This second phase should include involvement in standards implementation and establishment of the organization, management and engineering for the interagency internet. The last phase will emphasize the coordinated enhancement of functionality, capability and techn-ology for the interagency internet.

Standards Definition:

The Internet Organization will be charged to participate, on its own behalf and on behalf of the constituent networks, in national and international standards activities, to ensure that the international standards, and the networking products developed, adequately reflect the needs of the U.S. research community. At the present time these activities are dominated by the interna-tional PTT activities and by the industrial and computer manufac-turing companies, and the Internet (IP) and end-to-end transport (TP4) protocols, corresponding to the DARPA TCP/IP protocols, may not have the attention they require.

Interagency Research Internet Funding Requirements

Creation of an Interagency Internet will involve activities, such as hardware implementation, software development, communication lines/facilities, etc. which require funding over and above current agency budgets. This is true because current budgets have been developed and justified based upon individual agency research network requirements and not on interagency requirements. In addition, research into internetworking issues such as user authentication, access control, internetwork accounting, and internetwork performance will be undertaken by the IRI. There will also be a need to support administrative costs, consultants' studies, seminars, workshops, etc. for the IRI and these funds should also be included in the constituent agency budgets.

It is recommended that the necessary funding to accomplish this Interagency Research Internet capability be provided as additional funding to the constituent agencies as soon as possible so that this very important issue of interconnectivity can be addressed.

The resources required to create and operate an Interagency Internet will build up slowly as the level of activities grow. Annual funding in the out years is estimated to be approximately $7-$10 million per year. This is a relatively modest level of additional funding compared with the existing agency investment in supercomputers and supercomputer networking.

The First Step: An Interagency Network Task Force

To initiate the development of this Internet, it is recommended that a formal but temporary Interagency Task Force be established, comprising members from the Agencies who currently fund and sponsor networks to support research activities (DARPA, DOE, NASA, NSF). Other agencies may participate if required. This Task Force should be charged to codify and document the interagency research inter-networking requirements, to establish the detailed organizational structure for managing the Interagency Internet, and to request additional funding within the individual involved Federal agencies for implementing this interagency network.

Due to the importance of the interagency internet, this Interagency Network Task Force should be established at a sufficiently high level to command the resources necessary to implement recommendations. It is recommended that the Interagency Network Task Force be established immediately, and be charged with taking whatever steps are necessary to make rapid progress towards the definition of the Interagency Network requirements, and of the technical and organizational approaches to be taken. This Task Force should report to the FCCSET VHPC Committee by September 1, 1986.

Reference 5

THE EFFECT OF DISTRIBUTED COMPUTING TECHNOLOGY ON WIDE AREA NETWORK CAPACITY REQUIREMENTS, 1987, HALL, JOHNSTON, ROSENBLUM

The Effect of Distributed Computing Technology on Wide Area Network Capacity Requirements

Dennis Hall, William Johnston, Mendel Rosenblum

Lawrence Berkeley Laboratory
February 1987

Abstract:

This report identifies a need to increase wide area network capacity by as much as three orders of magnitude over the next ten years. These increases are necessary to support new distributed computing products. Such products increase productivity, but are currently available only on local area networks. There is no technical reason for limiting these products to tightly constrained geographical areas, however. They can operate perfectly well over any terrestrial distance provided sufficient bandwidth is available. Such bandwidth is available today with fiber optics. To quantify capacity requirements, network traffic generated by this newer technology is compared with traditional traffic in a local network environment. An extrapolation to wide area networks is made. Speculation about the long term future of distributed computing technology and its effect on network capacity requirements is offered. It is argued that an increase of network capacity by one order of magnitude is sufficient to accommodate new distributed computing technology on existing wide area networks. Two orders of magnitude are needed to accommodate a fully integrated distributed system with today's hardware speeds. Three orders of magnitude are needed to accommodate increases in hardware speed anticipated in the next five to ten years. Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.

Background:

Picture a scientist using a modern, high-performance workstation. One workstation window is opened to a supercomputer located fifty miles away. The supercomputer has been programmed to compute trajectories for a beam of heavy ions from accelerator description parameters stored in a file. The parameter file is displayed on a second workstation window.

Every few seconds, the supercomputer sends a thousand position vectors and a thousand momentum vectors to the workstation. The scientist notices that at each successive step, the momentum envelope is gradually expanding, yet the spatial envelope is holding constant. This indicates heating, and means the particles cannot be focused accurately on their target.

The scientist now opens another workstation window. An interactive program is invoked that allows the beam descriptor parameters to be adjusted while satisfying physical constraints expressed as differential equations. A new value for the field gradient in the focusing magnets is established, and the trajectory computation is restarted. The program halts immediately with the message: 'Similar values were tried previously: Run 145.' The scientist reviews the results of run 145, which have been conveniently displayed by the program, then tries an entirely new value for the magnet setting. The process restarts successfully. This time there is no indication of heating, and the simulated particles proceed to a tightly focused target area.

By using a computer model, an accelerator design flaw has been detected early, and an indication of the cure has been found. The scientist now redirects the graphic display files to a video recorder so that the results can be studied by the full design team at tomorrow's review. Members of the team are reminded of the upcoming meeting via electronic mail. A few of the images are selected to remind the scientist of the main simulation results. These are printed on a high quality laser printer located a few steps away, and shared by members of the design team.

The preceding scenario is possible using traditional distributed computing technology widely available today. However, these older facilities are more difficult to use and more likely to cause errors than newer products now emerging in the marketplace. These newer products currently operate only in the high speed, low latency environments of local area networks. But, their utility is by no means limited to geographically constrained environments. On the contrary, these newer services are useful over any terrestrial distance, if sufficiently high bandwidth is available. In the next section we review both traditional and more modern distributed computing products.

Distributed computing products - a review:

The following distributed computing facilities have been in common use for at least fifteen years, and could all be used in the above scenario. These are the so called *traditional* distributed services.

- *virtual terminal:* Virtual terminal facilities (e.g. [1]) provide interactive access to remote machines through a computer network. In the example, a virtual terminal facility could be used to access the remote supercomputer.
- *file transfer:* File transfer (e.g. [2]) allows files to be sent from one machine to another. In the example, file transfer could be used to send the particle positions and momentum vectors to the video recorder. It could also be used to send the parameter file from the workstation to the supercomputer and vice versa.
- *remote job entry:* Remote job entry (e.g. [3]) allows batch jobs to be submitted to remote machines. In the example, remote job entry could be used instead of virtual terminal for controlling the remote supercomputer.
- *electronic mail:* Electronic mail (e.g. [4]) is the computer analog of ordinary mail. It allows the exchange of electronic "letters." In the example, it could be used to remind the members of the design team of the forthcoming design review.

While the above services are extremely useful, most of them create an unnecessary interface layer between the user and the remote resource. The paradigm of these older services is to provide a remote service through a visible network access mechanism. These were designed in the days of low performance networks, when network bandwidth was scarce. By making the access mechanism visible, users were made aware of the resource they were consuming. This of course creates extra work, provides opportunities for mistakes, and causes a certain amount of frustration. It has the undesirable effect of reducing productivity.

High speed local networks have abundant network bandwidth. As a result, more fully integrated distributed services have emerged. In these services, the network is invisible. The following are a sample:

- *distributed printing:* The emergence of high speed local area networks along with low cost laser printers has dramatically changed the way printing gets done (e.g. [5]). It is now economically feasible to allocate high quality printers to a relatively small group, and to locate these printers in the user's work area. By attaching these to a local network, output can be routed to the appropriate printer regardless of where the output generating machine might be. Users merely request printing service, the routing is invisible.
- *network file system:* A network file system (e.g. [6]) makes files available uniformly throughout the network. The machine on which a file resides has no special status. In our example, a network file system could be used as an alternative to shipping the parameter file back and forth between the workstation and the supercomputer. Not only is the time and trouble of shipping files back and forth saved, but more importantly, errors that arise from inadvertently using an outdated copy of the file are eliminated. Users no longer need to be aware of where their files reside. They merely access files in the usual way, location is invisible.
- *remote procedure call:* A remote procedure call (e.g. [7]) is just the ability to call a procedure (or subroutine) on a remote machine. In the above scenario rpc's might have been used to split the trajectory computation between the workstation and the supercomputer. For example, the program running on the workstation might call several compute intensive subroutines on the supercomputer, which in turn might call graphics subroutines running

back on the workstation.

- *distributed window systems:* Windows provide a point for interaction between user and machine. They increase productivity by allowing users to perform tasks in parallel. A distributed window system (e.g. [8]) permits user level programs to perform complex graphical displays on another machine's window system efficiently. It decouples details such as scaling and positioning from the generating program. Graphical data may easily saturate even the highest capacity networks. Distributed window systems attack this problem by using a high level graphics description language (e.g. [13]) for communication. This reduces network bandwidth requirements while increasing graphical display functionality. The effect is to make available high speed, high quality graphics on machines (supercomputers for example) whose graphical support system is rudimentary compared to modern workstations. While this technology is still in its infancy, it already promises a major breakthrough in the way scientists interact with supercomputers.

The thesis of this paper is that demand for these newer distributed computing services on wide area networks will increase over the next five years. In the next section we discuss our reasons for this conclusion.

Why users need local area network technology on wide area networks:

Modern, distributed computing products increase productivity. Users accustomed to these products in their local environment will want them in their extended network environment. To illustrate this we compare a traditional implementation of our supercomputer scenario with one using modern distributed computing facilities. We emphasize that both approaches are fully implementable today using off the shelf technology. We assume our scientist is using a modern workstation in either case.

In a traditional network environment, our scientist might begin by opening a graphics device emulation window on the workstation. This window would be used just like an ordinary graphics terminal attached directly to the supercomputer. Within the graphics window, a virtual terminal utility would be invoked to access the remote supercomputer. The user would "login" to the supercomputer by providing ID and password information. The remote computation would then be started by providing the names of the program and its data files in a syntax acceptable to the remote supercomputer. The number of steps required to do this would be four: open graphics window, invoke virtual terminal utility, login to supercomputer, start program.

In a modern distributed computing facility, our scientist would open an ordinary workstation window, not a special graphics window, and would immediately start the remote computation by providing the names of the program and its data files in a syntax native to the workstation. No special graphics setup, no login procedure or other conscious access to the network would be necessary. Moreover, a single syntactic framework, that of the workstation, would be used throughout. Number of steps, two: open window, start program.

Note that no login step is required in the newer environment. This is normal in local area network environment where all users work for the same institution and security is not considered a problem. For remote supercomputers, security might well be a problem. To handle this the supercomputer could request a password whenever the time between accesses exceeds a threshold. The threshold could be chosen to allow users to work unhindered so long as a reasonable degree of interaction is going on. In the above scenario, our scientist would be prompted for a password at most once.

Returning to our comparison, when the scientist notices the beam is heating, the supercomputer computation is stopped in both cases. The next step is to adjust the parameter file. To adjust the parameter file in a traditional environment, our scientist first opens a new workstation window. The parameter file is retrieved from the supercomputer by invoking a file transfer utility. This of course requires providing a login sequence (ID and password) to the supercomputer. The parameter adjustment program is then invoked. Once a new value for the field gradient is established, the scientist returns the file to the supercomputer. The file transfer utility must be reinvoked and the login sequence must be repeated before the file can be returned to the supercomputer. Number of steps, five: login, retrieve file, adjust parameters, login, return file.

To adjust the parameter file in a modern environment, our scientist simply invokes the parameter adjustment program in the same window used to run the supercomputer computation. Number of steps, one: adjust parameters.

If the parameter file is large, and if the network is typical of traditional wide area networks (i.e. 56 kbit/sec land lines), transfer time could increase frustration. However, this effect is independent of the utilities used. It is an argument in favor of high speed wide area networks regardless of the sophistication of facilities.

The next step is to restart the supercomputer computation. Except for syntax, this step is the same in both environments. The program halts immediately with the message: "Similar values were tried previously: Run 145." At this point, the scientist using the traditional environment feels the first real pangs of frustration. Four of the previous five steps must be repeated. (The file needn't be retrieved from the supercomputer initially because it hasn't been changed.) The scientist in the modern environment only repeats one step, the actual parameter adjustment.

This time the program completes successfully. To create the video display in a traditional environment, the scientist must first retrieve the output from the supercomputer then direct it (as a local file) to the video recorder (two steps). In a modern environment the scientist simply directs the output file to the video recorder (one step). Sending mail to the design team, and selecting frames for printing are done the same way in both environments (two steps).

The scientist using the traditional network environment has performed eighteen steps while the scientist using the modern environment has performed seven. The effect on productivity is obvious. Run setup time is reduced or eliminated because all resources (remote and local) are accessed uniformly. Errors are less frequent because there are fewer opportunities for their occurrence and because a single command syntax is used. Frustration levels are lower because less time is spent waiting for results, and because low level tasks such as shuttling files back and forth between machines have been automated. As scientists become accustomed to these modern facilities in their local environments, demand for similar facilities in wide area networks will increase.

The modern distributed computing environment we have described above is in effect a single, integrated, nationwide "supercomputer." It would be accessed uniformly from anywhere on the network. Its total power would be enormous. Such a facility, available to the national scientific community, would create a technological and scientific environment superior to that of any country in the world. It would help to maintain the nation's competitive advantage in computing for decades to come. In the next section we quantify the effects of creating such a distributed computing environment on wide area networks.

A comparison of traditional and modern distributed computing traffic

The central theme of this paper is that significant increases in network capacity are needed if local area network technology is to be extended to wide area networks. Our experience in adding such facilities to the LBL local area network is outlined in the appendix. We feel this experience provides a forecast of what might occur if modern distributed computing services were extended to wide area networks.

We observe that network file system traffic per host is about an order of magnitude higher than traditional traffic. This is based on observations of diskless workstation traffic. Diskless workstations represent file traffic that can be expected in wide area networks when users must access files from more than one machine (as in the case of our scenario). Therefore, a wide area network that operates comfortably at 56 kbits/second might need a megabit per second (i.e. a T1 channel) to support a network file system or other modern protocols that function at this level of the operating system.

We further observe that more highly integrated services such as that represented by Sun's memory swapping protocol for diskless workstations *(network disk)*, create an order of magnitude higher load than the network file system. Although network disk would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks, it provides a tightly coupled service at a deep level of the operating system. As such it forecasts future distributed computing traffic on local area networks. We conclude that wide area networks operating at T1 speeds might need a 10 megabit fiber optic link to support traffic from future highly integrated distributed computing services with performance characteristics similar to Sun's network disk protocol.

So far our analysis has been based on performance of existing distributed computing facilities on existing hardware. The future will certainly bring increases in hardware speeds as well as more highly integrated network software. We think it is reasonable to project a factor of two increase every three years in available cpu power for the next ten years. Therefore, the wide area network load can be expected to increase another order of magnitude in ten years just from faster hardware.

In all, we project an increase of three orders of magnitude in wide area network capacity requirements. In other words, we think that in ten years scientists could use 100 megabit links from coast to coast to access a vast array of national scientific computers as a single, integrated "supercomputer." This would significantly increase the competitive edge of the United States in science and computing.

Conclusions:

(1) Productivity can be significantly increased by extending modern distributed computing facilities to existing wide area networks.

(2) As these facilities become commonplace in local area networks, demand for equivalent services in wide area networks will develop.

(3) To accommodate today's network file systems and other highly integrated distributed computing products, a factor of ten increase in network capacity is needed.

(4) To add software products anticipated for two to five years from now, an increase of another order of magnitude is projected.

(5) To assimilate hardware speeds expected in five to ten years, an increase of yet another order of magnitude is forecast.

(6) In all, an increase of three orders of magnitude in wide area network capacity requirements are projected for the next ten years.

(7) Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.

Appendix:

Experience with distributed computing traffic at LBL

The LBL local area network consists of a single logical ethernet spanning about half the physical area of the site (several square kilometers). To isolate and minimize traffic, the ethernet is physically divided into six segments. These are joined by bridging devices with address filtering. This confines network traffic with sources and destinations on the same segment, to that segment. A single probe, therefore, can only see network traffic on one such segment. This survey is limited to statistics on two of these segments: CSRLAN, the Computer Science Research segment, and CSLAN, the Computing Services segment. It is further restricted to TCP/IP traffic only. DECNET traffic, the dominant traffic on these networks, is not examined because DECNET currently provides only traditional distributed computing services on ethernets.

The two parameters used to characterize network load are packet rate and data rate. Tables 1 through 4 summarize the load on CSLAN and CSRLAN in time intervals ranging from 21 hours to 52 hours. Network traffic is summarized for all the traditional protocols as described in the preceding section. We have included distributed printing in the traditional services because today's wide area networks carry printing traffic, although often it is disguised as file transfer traffic. The left side of the tables characterize traffic as seen by the network. The right side characterizes traffic as seen by an "average" host.

Two lines in the tables are not described in the preceding section. The miscellaneous category covers 18 relatively uninteresting protocols ranging from the *internet control message protocol* to the *time* protocol for synchronizing host clocks. The user protocol collects a variety of user developed protocols. Some use Sun's remote procedure call facility. We expect such use to increase as remote procedure calls become easier to use and more widely available.

Key to protocol abbreviations	
vt	virtual terminal
ft	file transfer
ml	electronic mail
rje	remote job entry
pr	distributed printing
usr	user defined protocols
misc	miscellaneous
trad	all the above (vt-misc)
nfs	network file system
nd	network disk

Table 1: Traditional traffic on CSRLAN - 51.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	124	10.0	0.8	50.2	23.6	161	12
ft	59	0.6	0.2	2.8	7.6	19	8
ml	73	0.5	0.1	2.7	2.4	15	2
rje	18	1.4	1.0	7.1	30.0	157	108
pr	18	0.7	0.3	3.6	8.0	78	28
usr	17	4.0	0.6	20.3	17.5	475	67
misc	106	2.7	0.4	13.3	10.9	50	6
all trad	238	20.0	3.3	100.0	100.0	167	27

Table 2: Traditional traffic on CSRLAN - 27.52 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	93	9.2	0.7	37.0	13.0	197	14
ft	25	0.9	0.4	3.8	7.2	74	30
ml	58	0.7	0.1	2.7	1.6	23	2
rje	14	6.0	3.0	24.1	56.5	854	431
pr	17	1.0	0.3	3.9	5.7	112	35
usr	22	4.2	0.5	17.1	9.9	385	47
misc	70	2.9	0.3	11.5	6.1	81	9
all trad	161	24.8	5.3	100.0	100.0	308	66

Table 3: Traditional traffic on CSLAN - 41.99 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	79	8.0	0.8	30.0	14.9	203	19
ft	25	0.6	0.2	2.2	3.8	46	15
ml	54	0.4	0.1	1.5	1.0	15	1
rje	22	4.0	1.8	14.9	35.5	364	165
pr	32	9.3	1.7	34.6	34.0	581	108
usr	12	2.7	0.3	9.9	6.3	444	53
misc	66	1.8	0.2	6.8	4.4	55	6
all trad	154	26.9	5.1	100.0	100.0	348	66

Table 4: Traditional traffic on CSLAN - 20.74 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	38	6.8	0.5	13.9	6.1	360	28
ft	13	0.1	0.0	0.2	0.4	14	4
ml	39	0.4	0.1	0.9	0.6	22	2
rje	14	11.4	5.7	23.2	63.1	1630	807
pr	29	27.5	2.4	55.8	26.2	1895	162
usr	11	1.1	0.1	2.2	1.5	196	24
misc	51	1.9	0.2	3.8	2.2	73	7
all trad	100	49.2	9.0	100.0	100.0	984	179

Table 4 shows a significantly higher value of remote job entry and distributed printing load. These services are the two most heavily used in the traditional set, and use is increasing. Remote job entry traffic derives from the familiar remote shell [9] command in the Unix [10] environment. It allows commands to be executed on remote machines and files to be copied across the network. This latter use is functionally equivalent to file transfer except that no login is required. In the LBL environment, most remote job entry traffic is generated by disk backup daemons that wake up in the middle of the night to copy Unix disks onto the central VMS file cluster. The distributed printing service is also popular. It supports the Laboratory's distributed computing service that currently produces about a hundred thousand pages per month of printed output. The percentage columns show relative network traffic on the segment.

The packet and data rates for "average" hosts may be used to estimate the increase in network traffic that would be brought about by adding a protocol to a machine. Note that since all traffic is between two or more hosts, we double the network values to compute hourly rates per host. Clearly, remote job entry and distributed printing have the most significant effect in the traditional set.

The last line summarizes traditional network traffic. The network sees 20-50 thousand packets per hour and 3-9 megabytes of data per hour around the clock. An average host sees up to a thousand packets per hour and 30-180 kilobytes of data per hour around the clock. Note that the average traffic rates are much smaller than the rates for some individual protocols. This is because the denominator in the equation the number of hosts, includes all machines in the sample. Thus, average traffic is highly biased toward protocols that run on the greatest number of systems. We next show the effect of adding modern distributed services to this environment.

Tables 5 through 8 show traditional network traffic together with two modern distributed services: network file system and network disk. Network file system is the network file system described above, and network disk is Sun's network disk protocol (proprietary). Network disk provides a memory swapping service for diskless workstations. It would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks. However, it provides a tightly coupled service at a deep level of the operating system. We expect to see more such highly integrated services emerging in local area networks over the next two to five years. Systems such as Andrew [11] and Mach [12] provide a preview of things to come. Therefore, we feel network disk provides a good predictor for future services that will evolve first in local area networks and then be desired in wide area networks.

Table 5: Modern traffic on CSRLAN - 51.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	238	20.0	3.3	30.9	8.0	167	27
nfs	19	5.0	2.7	7.8	6.6	529	285
nd	16	39.6	34.9	61.3	85.4	4952	4360
total	238	64.6	40.9	100.0	100.0	542	343

Table 6: Modern traffic on CSRLAN - 27.52 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	161	24.8	5.3	43.9	16.5	308	66
nfs	18	4.7	2.8	8.3	8.5	521	307
nd	15	27.0	24.3	47.8	75.0	3598	3244
total	161	56.5	32.4	100.0	100.0	702	403

Table 7: Modern traffic on CSLAN - 41.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	154	26.9	5.1	68.1	32.5	348	66
nfs	8	1.0	0.4	2.5	2.5	245	99
nd	9	11.6	10.2	29.4	64.9	2575	2271
total	154	39.4	15.7	100.0	100.0	512	204

Table 8: Modern traffic on CSLAN - 20.74 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	100	49.2	9.0	79.9	43.9	984	179
nfs	10	1.6	1.0	2.6	5.0	324	202
nd	8	10.8	10.4	17.5	51.2	2691	2612
total	100	61.6	20.4	100.0	100.0	1232	408

On CSRLAN network file system traffic per host is clearly about an order of magnitude higher than traditional traffic, and network disk traffic is an order of magnitude higher than network file system. On CSLAN, this doesn't show as clearly because the workstations on CSLAN are primarily used as front ends to VMS machines, and because of the previously mentioned heavy use of remote job entry and distributed printing. Nevertheless, the network disk traffic is a factor of 20 higher than all traditional traffic. We conclude that a wide area network that operates comfortably at 56 kbits/second might need a megabit per second (i.e. a T1 channel) to support a network file system. Similarly, wide area networks operating at T1 speeds might need a

10 megabit fiber optic link to support traffic from future highly-integrated distributed computing services with performance characteristics similar to Sun's network disk protocol.

References:

[1] "Telnet Protocols," *The Arpanet Protocol Handbook,* Defense Communications Agency, Washington, D. C. 20305, Jan 1978, pages 89-264.

[2] "File Transfer Protocols and Standards," *ibid,* pages 265-332.

[3] "Remote Job Entry Protocols," *ibid,* pages 379-420.

[4] "Mail Protocol and Standards," *ibid,* pages 333-379.

[5] Johnston, W. E., and Hall, D. E., "Unix Based Distributed Printing in a Diverse Environment," *Summer 1986 USENIX Conference,* Atlanta, Georgia, June 9-13, 1986.

[6] NFS Protocol Specification, *Networking on the Sun,* Part No: 800-1324-03, Revision B of 17 February 1986. Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, CA 94043.

[7] RPC Protocol Specification, *ibid.*

[8] *NeWS Preliminary Technical Overview,* Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, CA 94043, October 1986.

[9] rsh-remote shell, *Unix Programmer's Manual,* 4.2 Berkeley Software Distribution, Virtual VAX-11 Version, March, 1984.

[10] Ritchie, D.M., and K. Thompson, "The UNIX Time Sharing System," *CACM,* Vol 17, No 7, July 1974.

[11] Morris, James H., et al, "Andrew: A Distributed Personal Computing Environment." *CACM,* Vol. 29, No. 3, March 1986.

[12] Rashid, R., "Threads of a New System," *Unix Review,* August 1986.

[13] Warnock, J. and Wyatt, D.K., "A Device Independent Graphics Imaging Model for Use with Raster Devices," *Computer Graphics,* Vol. 16, No. 3, July 1982.

Reference 6

FUNCTIONAL CHARACTERISTICS OF ESNET, 1987, DUNNING, GREENWALD, LOKEN

Functional Characteristics of ESNet: A Report to the Energy Sciences Network Steering Committee

T. H. Dunning, Jr., M. Greenwald, and S. Loken

The functionality provided by present-day local area networks (LANs), in addition to that provided by existing wide area networks (WANs), can be considered an adequate zero-order model for the Energy Sciences Network (ESNet). Ethernet-based LANS, with a bandwidth of 10 Mbps, are in operation in many research facilities and experience has shown that this type of network maximizes the utilization of both computer and human resources. The WANs utilized in the energy research community are mostly based on DECNET, MFEnet and TCP/IP protocols.

The desired features of ESNet are summarized below. As user needs vary significantly and networking costs are strongly dependent on the level of service provided, ESNet will need to offer different classes of service, e.g., some users/sites may only be able to support 9.6 kbaud service while other users/sites may well require T1 (1.5 Mbaud) service. The service level will be determined by (among other factors) the size/budget of the project, the scientific importance of the project, and the importance of the capabilities provided by ESNet to the success of the project.

- Network Control and Management

The network management system should allow the addition and deletion of nodes locally, with an automatic update of the necessary routing tables. The total number of nodes allowed should be essentially unlimited. There should be automatic rerouting when a path fails. Network trouble should be reported to the affected users as soon as possible. The network manager should be able to logically disconnect a node.

- Terminal/Workstation Access

Terminal/workstation access must be provided in a transparent manner. Network latency must not impede terminal/workstation access to remote computers; this may well eliminate the use of satellite communications links for interactive work and places constraints on the performance of the network.

Full screen editing should be supported. All user data should be transmitted without modification. Network oriented commands (such as NETOUT and NETPLOT) should be available in two flavors: the first set should be adapted to the specific user environment (e.g., VMS, CTSS and UNIX), while the second set should provide a uniform interface in all of the major user environments.

The network should be treated as a scarce resource. Users should be encouraged to make efficient use of the network. The managers of ESNet should encourage the computer and operating system vendors to develop networking products which provide services such as shadow editing, automatic compression/decompression, etc. Further, to take advantage of the rapidly increasing capabilities and decreasing cost of scientific workstations, integrated workstation/network/supercomputer software should be developed to make effective use of both computer and human resources.

For both interactive and batch work on remote computers, users require adequate communication speeds. Experience on local computer networks indicates that communication speeds equivalent to 9.6 kbaud or higher is required to effectively utilize both machine and human resources. Further, for applications requiring heavy use of graphics, communication speeds equivalent to 56 kbaud or higher will be required. This will require T1 (1.5 Mbaud) service for all major sites within the near future.

- Electronic Mail Services

Electronic mail provides an effective means of communicating with other scientists or groups of scientists in an expeditious manner. Electronic mail greatly facilitates collaborative efforts; such efforts greatly increase the effectiveness of central facilities. ESNet should fully support electronic mail services.

Electronic mail service should provide automatic logging of mail received, automatic notification of sender of receipt of mail, and automatic forwarding. The Mail system should be store-and-forward so that messages will be delivered eventually, even if the path is momentarily down. Mail forwarding should be capable of reaching all nodes on ESNet regardless of their connection path. Finally, an electronic address book should be maintained on-line; this will make it possible to send electronic mail without prior knowledge of the intended recipient's account name or node.

It should be possible to include any kind of file (not just text) in a mail message.

Bulletin board services should be supported for topics of broad appeal to the energy sciences community. Access should be provided for other bulletin board services of interest.

- File Transfer

The file transfer protocols should support all file types; ASCII and binary files should be transmitted intact. The file transfer system should be designed so that files will be delivered eventually, even if the network is momentarily down.

There should be both prioritization and interleaving of files so that long files do not unduly impede the flow of short files. A high speed, less time critical link should be used for file transfers. The user should be provided with the means to schedule file transfers for off-hours transmission.

Distributed network file systems should be supported as they become available.

The managers of ESNet should consider the establishment of a multicast capability based on low cost, small aperture satellite systems. The multicast will allow very high speed data transfer, which could be large files, graphics, or other data, to many locations simultaneously. This capability could be used to send a graphics control system display to many collaborators during an important "event". It could enable software or information dissemination throughout a collaboration for "current releases" or other changes. Further, satellite technology is the least expensive way to get high bandwidth communication to remote and/or smaller sites which are either not served by high speed switches in their local geographic area or which have a small number of users and cannot justify more expensive facilities. Finally, large scale computational simulations often require only small amounts of data in but lots of data out. Small VSATs can be installed in a smaller, cheaper receive-only configuration which provide high speed in and which complement a lower medium speed out. Thereby providing enhanced functionality without the need to install full terrestrial T1 circuits.

ESNet should provide the highest data transfer rate possible consistent with budget and technology.

- Remote Printing/Plotting

The emergence of low cost laser printers/plotters has made it possible for small groups to have local access to high quality printing/plotting devices. The user should be able to direct files to any "supported" printer/plotter on any node on ESNet for printing/plotting. This may require the adoption of a standard metafile format and/or standardization on a few printers/plotters. An initial adaptation of postscript is encouraged.

- Interprocess Communication

Support for interprocess communication is necessary for implementing the other requirements, e.g., for the efficient implementation of full screen editing, workstation support, etc. While the development of these capabilities is beyond the scope of the network per se, the design of the network must not preclude nor impede these applications. This capability is also likely to be important for distributed CAD/CAM applications as well as the research activities envisioned here.

. Interconnection to Other Networks

ESNet must be committed to supporting international networking standards and to the timely migration of ESNet to the ISO standard. This should greatly facilitate interconnection to other networks.

ESNet must provide interconnections to local areas networks (and clusters) attached to the ESNet nodes. ESNet should be connected to the networks being developed by other science-oriented government agencies, specifically, NSF and NASA. Transparent access should be provided to these networks. Gateways should be provided to other networks used extensively by scientists, e.g., BITNET.

. Distributed Window Systems

ESNet should support distributed windowing systems. Windows increase productivity by allowing users to perform tasks in parallel. By using a high level graphics description language, distributed window systems can substantially reduced the load on the network resulting from graphics applications; this reduced bandwidth requirements while increasing graphical display functionality.

. Remote Job Entry

Job entry should be supported on any computer system on the network for which the user has access.

. Network File System

The network file system should make files available uniformly throughout the network without regard to the physical location of the file so long as it is on a "supported" node, e.g., it should allow access to (reading) the contents of remote files without file transfer. This eliminates the wholesale transfer of files from one machine to another and eliminates the possibility of errors arising from the use of outdated files.

. Document Transfer

The network should permit the transfer of scientific documents. This is more than the transmission of text files. A standard Document Interchange Format (DIF) should be used to transfer formatted documents with merged text, graphics, images, and data.

FAX and freeze-frame video conferencing services should be supported.

Reference 7

HENP, FE, AMS, HER, AND BES PROGRAM REQUIREMENTS, 1987

HIGH ENERGY & NUCLEAR PHYSICS

NETWORK REQUIREMENTS

HIGH ENERGY PHYSICS COMPUTER NETWORKING

Response for COMPUTER NETWORKING STUDY requested by:

John S. Cavallini
ER-7 Scientific Computing Staff
Office of Energy Research.

Prepared by Networking Representatives of:

Argonne National Laboratory
Brookhaven National Laboratory
Fermi National Accelerator Laboratory
Lawrence Berkeley Laboratory
Stanford Linear Accelerator Center
L3 Collaboration

Transmitted by:

H.E.Montgomery
Fermi National Accelerator Laboratory

DECEMBER 19 1986

I. INTRODUCTION

High Energy Physics (HEP), also known as Elementary Particle Physics is a branch of physics research which explores the underlying physics of the universe. It employs techniques involving large accelerators which permit the investigation of matter on a minute, sub-nuclear, sub-nucleon scale. The energy scale involved is such that the only natural occurrence of such energies currently known is that which existed in the very early universe, a fraction of a second after the Big Bang. This links the basic ideas of HEP rather intimately to Cosmology.

The size of experiments, the scale of accelerators and the cost of the efforts have led, over the last 40 years, to the concentration of the experimental sites at relatively few (~10) laboratories worldwide. The majority of the practicing High Energy Physicists are however University based. The typical experimental collaboration today encompasses from 100 to 400 physicists, who come from 10 - 40 institutions, in as many as a dozen countries.

Within the USA the primary funding sources for these efforts are the Division of High Energy Physics within the DOE and the Physics Branch of the National Science Foundation. The DOE has been designated by the Congress as the "Executive Agent" for the US high energy physics program.

The DOE and the NSF are advised as to the direction of the program by the High Energy Physics Advisory Panel (HEPAP) made up of prestigious members of the HEP community.

II. HISTORICAL NOTES

In 1984 a sub-panel of HEPAP was formed under the leadership of J. Ballam of SLAC, to consider the Computer Needs of High Energy Physics for the Next Decade. The sub-panel produced a report entitled COMPUTING FOR PARTICLE PHYSICS (DOE/ER-0234) in August 1985. The report included an Analysis of Networking Needs, prepared by sub-panel members and associated working groups of specialists, which is attached to this reply.

During the latter part of 1985 a group of interested members of the community responded to the recommendations of that paper and produced a draft document entitled 'Prospectus for a High Energy Physics (HEP) Network'. The Prospectus was included by some of the laboratories in their responses to a questionnaire on the needs for Energy Research Networking in late 1985 - early 1986, which had been generated by the Office of Scientific Computing.

All of these documents have emphasized the mode of growth of the currently installed networking facilities, specifically that the current installations were put in place in response to immediate needs. Funds for the facilities came from general university and laboratory operating budgets, and therefore have not been recognized until now (with a few exceptions) as a separate activity supported by central funding.

In most cases, the wide range of networking services required led to nearly independent implementations for terminal access to the larger computing resources at the accelerator sites, and for computer to computer communications between home institutes and other collaborating institutes.

III. MODUS OPERANDUM

High Energy Physics is a strongly collaborative discipline. This is especially evident in the large collaborations of Experimental HEP, but the trend is also increasing in theoretical projects. The experiments are performed at large accelerators which also provide large computing installations for offline data analysis. The online local control of experiments is performed by several networked computer systems in the Super - Minicomputer class. At many laboratories the online and offline computer systems are linked to each other and to the outside world by site-wide networks.

Super - Minicomputers at the home institutions of each collaborating HEP group are used for the group's own computing needs, or often they are run as Departmental computers. The need for networking between the computers at the experiment, at the Computing Center in the accelerator laboratory, at the home institutions, and at centers with special computing resources, is a necessity for the functioning of the field. The pace of the research is such that it is impossible for the University Professor to limit his participation to weekend visits and summer stays at the experiment or with his colleagues. The computer network has to provide the physicists with the ability to:

- have terminal access to computing resources throughout the accelerator site. Whether at the experiment, in the computing center or elsewhere on the lab site, the access modes must be matched to the computing resource.

- use electronic mail implementations to communicate with colleagues across the world, immediately and in the precise way permitted by the written word.

- transfer files between computers at any two member sites in the collaboration, or at the lab site. This allows the physicists to exploit all of the computing resources available to the collaboration in an efficient manner. It also allows them to share in the development and use of the analysis and simulation software, and the databases of experimental parameters.

- submit computational tasks to any computer on which he has authorization.

- obtain the results of computations remotely, whether in graphic or textual form.

- utilize distributed data base techniques to facilitate the preservation of data integrity when physics analyses are distributed throughout the collaboration.

- access libraries of code which may be generated and validated centrally, but which are then distributed to all the home institutions.

- perform communication between tasks running on different computers. This facilitates a high degree of distributed processing, and enhances the efficiency of use of the global computing resources. It also permits the monitoring of running experiments from remote sites.

IV. EXISTING NETWORKS

The original growth of networking for HEP started several years ago (before 1980), at a time when there was little appreciation of the global impact of networking, or of the need for "standards". As summarized here, High Energy Physicists therefore used a variety of networking facilities to satisfy the ensemble of their requirements.

The division into categories used to construct the list in this section is designed to facilitate the reporting process. It should be understood that the different items in the list are rather different entities.

It should also be noted that the HEPAP sub-panel recommendation for HEPNET, and the Prospectus, did not specify a single networking protocol. It was recognised that a prerequisite of any implementation was the provision of the TOTAL FUNCTIONALITY offered by the ensemble of services listed below.

i) Data Switch Networks

Terminal access to the Central Computing facilities at the laboratory sites has been extended off - site to some of the collaborating physicists' home institutions. There is also a connection between the switches at FNAL and SLAC, which means that anyone accessing the switch at one site also has access to the switch at the other site. These networks utilize both leased lines and dial-up connections. Dial-up tends to be an initial implementation, which is rapidly superceded by more cost effective leased lines. In this section we also include the remote printing facilities examples of which are the TCF protocol associated with the Cyber system at FNAL and the RSCS protocol associated with the IBM at SLAC.

ii) BITNET and EARN

This network is one of the most extensive in the USA (BITNET) and in Europe (EARN). It also extends to Asia. The network is accessible to the general academic community, and for this reason it has been a natural vehicle for MAIL and light, non - urgent FILE TRANSFER traffic. It utilizes leased lines to 'nearest neighbors', and a store - and - forward file and message transfer technique.

iii) DECNET

The most common type of Minicomputer in use, both as a Departmental computing resource and in the HEP experiments in all laboratories, is the VAX manufactured by the Digital Equpment Corporation (DEC). As a result, networks using DECNET have proliferated between member institutions within collaborations, and as a major component of the Local Area Networks at each of the laboratories and at many universities. This ensemble of DECNET's has now merged, to a great extent, into a 1000 - node single network which covers the U.S. and part of Western Europe. The total number in Area Codes 41, 42 and 43, those primarily associated with H.E.P. in the U.S. is about 300.

iv) LEP3NET

A major component of European Networking strategy has been based on acceptable international standards for the lower network layers. This led to the use of X.25 as the unique standard used in the public packet switching networks operated by the national PTTs within each country.

There has therefore been a very strong motivation for collaborations which are based in Europe, or which have European institutions participating, to implement an X.25 network. A collaboration - wide network can then consist of both leased lines, and connections across one or more public networks in a transparent manner. The use of X.25 is also fully compatible with the simultaneous use of several higher level protocol suites, including DECNET, the COLOURED BOOKS, and the ISO standard protocols.

L3, which has an experiment under construction at the LEP accelerator at CERN, has therefore built the LEP3NET network based on X.25 switches. The network uses leased lines, with transparent gateways at several points to the world's public networks. LEP3NET also provides DECNET service, which includes an interim connection between the U.S. and European DECNET networks used by HEP.

v) Other Networks

Another network of interest to the High Energy Physics community is MFEnet, which currently is used for access to the Energy Research Supercomputing facilities at the MFECC and FSU, and to ARPANET.

HEP collaborations which include groups funded by NS are also considering the potential use of the NSFnet backbone for access to the NSF Supercomputer centers, to the ARPANET, and to a rapidly increasing number of university campuses.

The U.S. public X.25 networks, TELENET and TYMNET are also used by an increasing number of physicists. SLAC, LBL and FNAL have TYMNET connections, and both TELENET and TYMNET may be used to reach CERN. Access to LEP3NET over the public networks inside the U.S. may be used to reach CERN with reduced public network charges.

V. ORGANIZATION

The day to day operation of the network(s) is the responsibility of the local management at each of the nodes and access points. Coordination is offered through the HEPNET TECHNICAL COORDINATING COMMITTEE (HTCC). This committee was formed at the request of W. Wallenmeyer of the DHEP program office in DoE, and it reports to him through a representative of FNAL. The committee contains representatives of the HEP laboratories: ANL, BNL, FNAL, LBL, and SLAC, and of the L3 collaboration. Its meetings, which are held three times per year, are open to the community. Much of its business is conducted by network communications.

VI. CURRENT IMPLEMENTATION

As mentioned earlier Networking within High Energy Physics has been and is multifaceted. The comments in this section are to be taken in that context.

1) Network Topologies

i) Data Switch Networks

The Data Switch networks are based at the various accelerator sites which possess large central computing facilities. Typically, access is provided from the switch to a large variety and a large range of facilities throughout the site, including all major computing systems. Taken as a whole, the Data Switch network consists of a number of stars based at ANL, BNL, FNAL, LBL, and SLAC, and to a lesser extent at MIT. There are also a limited number of interconnection lines which permit intersite through-access, for example, from the Fermilab PBX to the SLAC PBX and vice versa.

ii) BITNET

BITNET, as stated earlier, has no imposed structure, and it has grown as a succession of links from existing network nodes to adjoining "nearest neighbors". The CUNY to Univ. California at Berkeley link serves as an East - West trunk line.

iii) DECNET

DECNET growth was originally based on the needs of individual collaborations for access to, and communication with, their Data Acquisition computers. Transcontinental connection was initially achieved when an institution such as the U. of Illinois at Champaign participated in experiments at both SLAC and FNAL. There are currently direct lines between BNL and FNAL, and between FNAL and LBL which provide reasonably direct connections across the continent. Virtual LEP3NET circuits between Caltech, Michigan, MIT and CERN also carry DECNET traffic across the country, and to Europe.

iv) LEP3NET

The LEP3NET network in the USA has a topology which is basically a tree structure with the primary star at MIT-LNS and secondary stars at Michigan, Caltech and CERN in Europe. At each of these sites there are X.25 switches. In addition there are individual nodes connected directly to the main switch at MIT.

2) Geographic Locations

The Geographic locations associated with the different elements of the networking functionality are shown in the successive figures VI.2.1.....

3) Protocols

i) Data Switch Network

ii) BITNET

BITNET is based on IBM's VNET and consists of store and forward computer to computer links. 9.6 kbps and 14.4 kbps modems drive point to point leased lines. The IBM RSCS protocols are used at the network and link layers and bisynchronous communication (BSC) at the link layer.

iii) DECNET

DECNET is the acronym for the Digital Equipment Corporation network offering. It is a full suite of protocols corresponding roughly but not exactly to an OSI layering. There are multiple possibilities at the data link level including Ethernet, DDCMP and X.25. DECNET permits rather transparent file transfer, and task to task communication including record level access to files, between two DEC computers at remote sites on the network. DEC are on record as to their commitment to a movement to ISO standards.

iv) X.25

X.25 is the name for the physical, link and network layer protocols used in many packet switching networks, as defined by a series of CCITT standards. X.25 is widely used by public packet switching networks in many countries. It offers a useful base for private networks since a wide range of higher level protocols are compatible with X.25. The latest version of the X.25 standard, which was released in 1984, permits up to 128 outstanding packets and therefore can be used efficiently over one or more satellite hops.

4) Installed Equipment,Communication Lines

The emphasis in this document is on Wide Area Networking. However, it should be recognised that a major part of the installed equipment at the laboratories and other sites is associated with local networks and services. The division between the two is difficult to make in general and differs from item to item. For example, on-site communication lines are excluded from the section on communication lines, but a description of the Data Switch systems at each lab is included.

i) Data Switches

a) ANL

ANL is in the process of installing an Intecom S/80 voice/data PBX system which will allow them to provide Data PBX services as well as Ethernet terminal service throughout the site. Currently X.25 switches are used for access to the IBM 3705 system. Some individual divisions use MICOM or Gandalf switches for access to Vax minicomputers.

b) BNL

BNL has a Gandalf Pacx 1000 system which can service up to 1024 terminals. This switch is networked to a smaller satellite switch used by MIS. The NSLS will acquire a small switch in the near future. With the extension of an Ethernet LAN all the major systems, including the IBM mainframe, will be accessible for remote login from any other local BNL system.

c) FNAL

FNAL has a distributed system of Micom Data PBX switches. There are 6 major and 4 minor systems connected together to provide service for 2500 terminal connections. There are 14 Universities connected to the FNAL system by leased lines and statistical multiplexors for remote access.

d) LBL

LBL has a Develcon Data PBX system serving approximately 2000 lines. In addition there is a Gandalf Data switch serving approximately 100 lines. The LBL ethernet is extensive and is the communications backbone among several hundred local computers.

e) SLAC

SLAC has a Micom Data PBX system which provides service for about 1500 terminal connections. SLAC also uses Bridge Communications Ethernet terminal servers (running XNS protocols) supporting about 150 terminals. There are 14 universities and laboratories (including MFECC) connected to the SLAC systemby leased lines and statistical multiplexers for remote access.

ii) BITNET Gateway

While many of the HEP sites are BITNET nodes there is also a mail gateway between BITNET and the HEP DECNET supported at, and by, LBL.

iii) DECNET Routers

Decnet routers are used as the primary means of DECNET distribution within the HEP DECNET. At Level 1 for local routing a number of machines are in use, at Level 2 and for inter-area routing there are a total of less than 10, one router at BNL, seven at FNAL and one at LBL. These along with several Vaxes provide the backbone function for the DECNET traffic.

iv) X.25 Switches

CAMTEC X.25 switches are installed at MIT, Michigan, Caltech and CERN. The MIT Switch has a capacity of up to 18 synchronous and 8 asynchronous incoming lines. 4 of the ports are capable of speeds of up to 64 Kbps, with a total switching capability of approximately 200 packets per second. The switches at Caltech and Michigan are limited to 19.2 Kbps and 40 packets per second, but may be upgraded in the near future. CERN has several switches, some of the type installed at MIT, which support connection to several European public and private networks in addition to LEP3NET.

v) Communication Lines

Table VI.4.v.1 contains a compendium of the leased lines used by the High Energy Physics community. The end points of each line, the bandwidth, the protocols supported and other information are also given. The cost of each line is of the order of $ 12,000 per annum for 9.6 kbps. Mail Gateways Between HEPNET and ARPAnet, and between HEPNET and MFEnet are installed and maintained at LBL.

VII. CURRENT REQUIREMENTS ASSESSMENT (4-5 Years)

In order to estimate the detailed requirements it is necessary to have a model for the expanded/enhanced network.

We therefore assume that in the short term future, the High Energy Physics network will be implemented as an X.25 network - level service running over leased lines, and supporting multiple higher - level protocol suites. This choice is driven by several factors:

1) An X.25 network service may be built today, starting with currently available commercial products, which will support the currently installed network capabilities. It is also an integral part of current MFEnet II plans.

2) Experience with X.25 private networks within the HEP community has demonstrated the functionality and reliability of this solution for a range of networks.

3) For international connections, X.25 is an essential part of the network service.

4) X.25 is part of the OSI standardand therefore naturally supports the evolution to the full set of OSI standards.

5) X.25 will support a multi-vendor environment.

Based on these considerations X.25 will satisfy the two primary goals of any implementation:

A. support of all the currently installed capabilities, and

B. evolution to future standards.

An analysis of the packet switching capabilities required, assuming the use of 56 Kbps X.25 inter - laboratory backbones, and many 9.6 Kbps feeder lines between each home site and a principal node at an HEP laboratory site, has been carried out by Howard Davies. His analysis is summarized in a SLAC Memo dated Sept. 19, 1986 which is attached as Appendix I.

The analysis was motivated by the need to determine whether relatively inexpensive present - day commercial devices could fulfill the packet switching needs. It was also necessary to set performance requirements for the proposed MFENET II gateways, which are intended to provide a gateway - to - gateway X.25 transport service across an IP based network.

The Davies study concluded (see Appendix I) that although the packet throughput requirements for the principal switching nodes are relatively high (200 packets per second), the needs could be met by existing commercial devices. Since the date of this study, new devices have come to the attention of the HTCC that reinforce his conclusions.

It is therefore assumed that the HEP network will progress through the following stages:

a) The network will develop a homogenous X.25 base, which will permit the utilization of multiple higher - level protocol suites. This will satisfy the requirement that at no stage will there be a reduction in the full range and functionality of the network services, relative to the facilities currently in use in HEP.

An X.25 base, up to the network layer, will allow LOGIN traffic, DECNET, Coloured Books and TCP/IP services to run without mutual interference. This will permit the efficient use of the bandwidth available on the backbone lines of the network, since each user may have to access which ever network community he requires, across the same set of lines.

b) If TCP/IP emerges as the interim de facto standard protocol of choice, the X.25 base will permit the introduction of its use to be accomplished naturally, without interrupting the full range of ongoing network services. The initial X.25 base will also permit the utilization of the MFEnet II Backbone connections between MFEnet II Gateways which are intended to support both X.25 and IP services across that network.

We note that TCP/IP has a significant multi-vendor base.

c) The migration path to ISO - standard protocol suites will be possible within the next 5 years. A large set of relatively stable ISO products from commercial vendors is expected by 1988. Full stability for most vendors' network products is possible in a time frame of 5 years. It should be noted that there is significant evidence that a range of important vendors are committed to the support of ISO standards, these include DEC and IBM. In particular development of FTAM and X.400 products is underway.

A step - by - step solution to the HEP networking problem could be implemented as follows:

FY87

— establishment of X.25 switches at MIT (existing), BNL, FNAL/ANL, LBL, SLAC and FSU. An upgrade of the MIT switching capability may be necessary.

— increase of the ANL/FNAL – LBL/SLAC backbone to 56 kbps.

— reconfiguration of the East Coast area network connections to permit cost effective utilization of a 56 kbps East Coast – ANL/FNAL backbone line by many Universities in New England and some Mid - Atlantic states. This reconfiguration should also permit a more efficient, generally accessible interconnection of the US DECNET to the L3 Transatlantic line. A proposal is in preparation from a group of these universities.

— connection of BNL and MIT by a 56 Kbps backbone. This would ensure adequate connection of all the US labs to the L3 Transatlantic line.

— establishment of a 56 kbps USA-CERN satellite link. This link would provide urgently needed capacity for a higher volume of file transfers than is now possible. This link is the subject of an existing proposal from FSU.

It must be noted however that current implementations of remote terminal service over satellite links have serious limitations. The existence of this link will not obviate the need for a transatlantic terrestrial line.

— establishment of a 56 kbps satellite link between the USA and Japan. This is the subject of a current proposal from three USA - Japan collaborations working at FNAL, BNL and KEK (Japan).

— installation of the first elements of a general network management system. This system will monitor the traffic flow pattern and its growth, localize and help trouble shoot network problems.

FY88

— installation of X.25 switching facilities at Cornell. Enhancement of the X.25 switches at BNL, ANL/FNAL and LBL/SLAC to the point where all the major laboratories have independent switches.

— establishment of 56 kbps line from Cornell to either BNL or MIT.

— improvement of end node connections such that appropriate lines and switching or multiplexing facilities are available with up to 9.6 kbps bandwidth for each service at all moderately sized University HEP Departments, with appropriate scaling for larger and smaller Departments.

— migration to the use of MFEnet II gateways, where feasible. Assuming that MFEnet II were fully implemented in this time frame, utilization of the MFE lines could be economically advantageous.

— introduction of TCP/IP services. This would allow increased utilization of high performance Graphics Workstations for local analysis of processed data from the Supercomputer or accelerator Computing Centers. We note that TCP/IP may offer improved multi-vendor capabilities, at least in the U.S., in the period before full ISO Standard OSI implementations become generally available.

— complete installation of the network management system.

FY89

— the only major topographical change that may appear is the addition of the SSC site as one of the principal switching nodes. This is contingent on SSC approval and site selection.

— the establishment of a major trunk between the East and West coasts. This would provide a necessary cross - country performance increase, and a higher level of reliability since the previous backbone links would provide some redundancy.

— upgrade of all trunk lines to at least 112 Kbps (2 X 56 Kbps). If current cost patterns persist, it will be most cost effective to combine satellite channels at speeds above 56 Kbps (for bulk file transfer), with terrestrial lines at 56 Kbps (for interactive and process - to - process applications requiring rapid response).

— installation of a 56 Kbps link between the U.S. and CERN over TAT-8, the fiber optic Transatlantic cable scheduled to start in late 1988.

— in general it is desirable that the services perceived across the network are equivalent to those provided on systems which are local to the user. It should therefore be foreseen that all feeder lines are upgraded to provide an effective speed of 19.2 kbps for all network services.

— begin migration to ISO products by the installation of ISO network interfaces and software products on a test basis at each laboratory site.

— begin migration to an ISO network management system.

Note: plans for FY90 - FY91 are contingent on the state of network technologies in this time frame. Key developments which are expected in this time frame, which will make the plans economically viable are:

(1) Widespread availability of fiber optic digital links throughout the U.S. Current trends indicate there will substantial excess capacity by the end of 1989.

(2) Commencement of TAT-8 in late 1988.

(3) As a result of (1) and (2) major cost reductions per unit bandwidth.

(4) The appearance of ISO products on a wide scale starting in 1988. Stabilization of the market, for all major computer vendors by 1990 or 1991.

(5) The appearance of inexpensive network interfaces and packet switching equipment capable of supporting higher data transmission speeds. This is expected as a natural outgrowth of 32 - bit microprocessor developments which are already underway.

FY90 - FY91

– upgrade of the link between CERN and the U.S. to at least 2 X 56 Kbps by FY90, as data taking at LEP begins. Progressive upgrade of this link in FY91 - FY92 if feasible.

– progressive upgrade of all major U.S. backbones to T-1 speeds (1.544 Mbps).

– progressive upgrade of all feeder lines in the U.S. to 56 Kbps.

– choice of satellite or fiber optic links, according to the evolution of technologies and prices, for bulk file transfer.

– experiments with higher speeds (above T-1) for bulk file transfer and data distribution if economically feasible.

– migration to ISO network protocols and associated hardware and software products.

VIII. SECURITY REQUIREMENTS

None of the work engaged in by High Energy Physics laboratories is classified, indeed one of the characteristics of the field is the indifference to international boundaries in the conduct of its work. On the other hand, there are increasingly frequent incidents of break-ins to computer systems, through remote login. These unauthorized individuals, by intent or ignorance, are a threat to data bases and files which represent large investments of intellectual effort and time. The level of security should therefore be sufficient to discourage such intrusions while providing minimal inhibition to the efforts of the bona fide practitioner.

X.25 has management capabilities which will permit a significant degree of access control, and which may deny the opportunity for attempted breakins by to those who are not located at recognized sites matched to authorization codes.

IX. USER INTERFACES and SPECIAL REQUIREMENTS

The choice of user interfaces and special devices, indeed the computers themselves, will be the concern of the individual laboratories and universities. It is intended that the network support be sufficiently broad that it can support a large fraction of the available devices. This is achieved in the short term by the support of multiple protocols and in the longer term by migration to standards as they evolve.

X. NETWORKING RESEARCH NEEDS

There are no requirements for basic Network Research since the HEP community is principally concerned with the functionality that it obtains as a USER. It does not consider that its needs exceed the bounds of networking technology which is currently available, or which is likely to become available in the near future.

On the other hand, it must be recognized that the testing of newly available equipment, which is acquired to satisfy its networking requirements, demands resources in terms of manpower, test equipment, time and money. We assume that these are integrated into the management and support functions, and are adequately provided for.

A modest level of support will be needed in addition for the testing and implementation of network management tools.

The widespread adoption of ISO network standards in a rapid but organized fashion, when the necessary products first become available, will also be an activity which is properly considered as "application level" research.

XI. ANNUAL FUNDING REQUIREMENTS

The funding requirements summary given below is based on the model, and the progressive steps outlined in Section VII. The costs listed include partial support for connections and upgrades of connections of end nodes at Universities, where major network needs for DOE supported programs cannot be supported by the ongoing operating budgets. It is recognized, nonetheless, that priority will be given to the backbones and to the basic network infrastructure, if central funding levels are not sufficient to cover the end node costs listed below.

The manpower estimates given include only the additional staff required to support a general service for the HEP Community, over and above the ongoing network activities which are of local concern to each laboratory site. We assume that the local concerns will be supported, as in the past, by the existing network staffs.

	FY	87	88	89	90	91
Trunk Lines	($k)	175	225	265	325	375
International Links		200	200	250	250	300
Software	($k)	60	65	70	80	100
X.25 Switches ($k)		100	75	60	70	80
Network Managment Systems, Test Equipment	($k)	75	75	50	50	50
Support for End Lines	($k)	125	150	180	215	255
Modems, interfaces, multiplexing equipment	($k)	100	80	90	90	100
Travel	($k)	55	65	70	70	70
Total (not including salaries)	($k)	890	935	1035	1150	1330
Manpower	(FTE)	1.5	2	2	2.5	2.5

[Five year plan cost $ 5.34 M + salaries .]

HEP LEASED LINES FOR TERMINALS

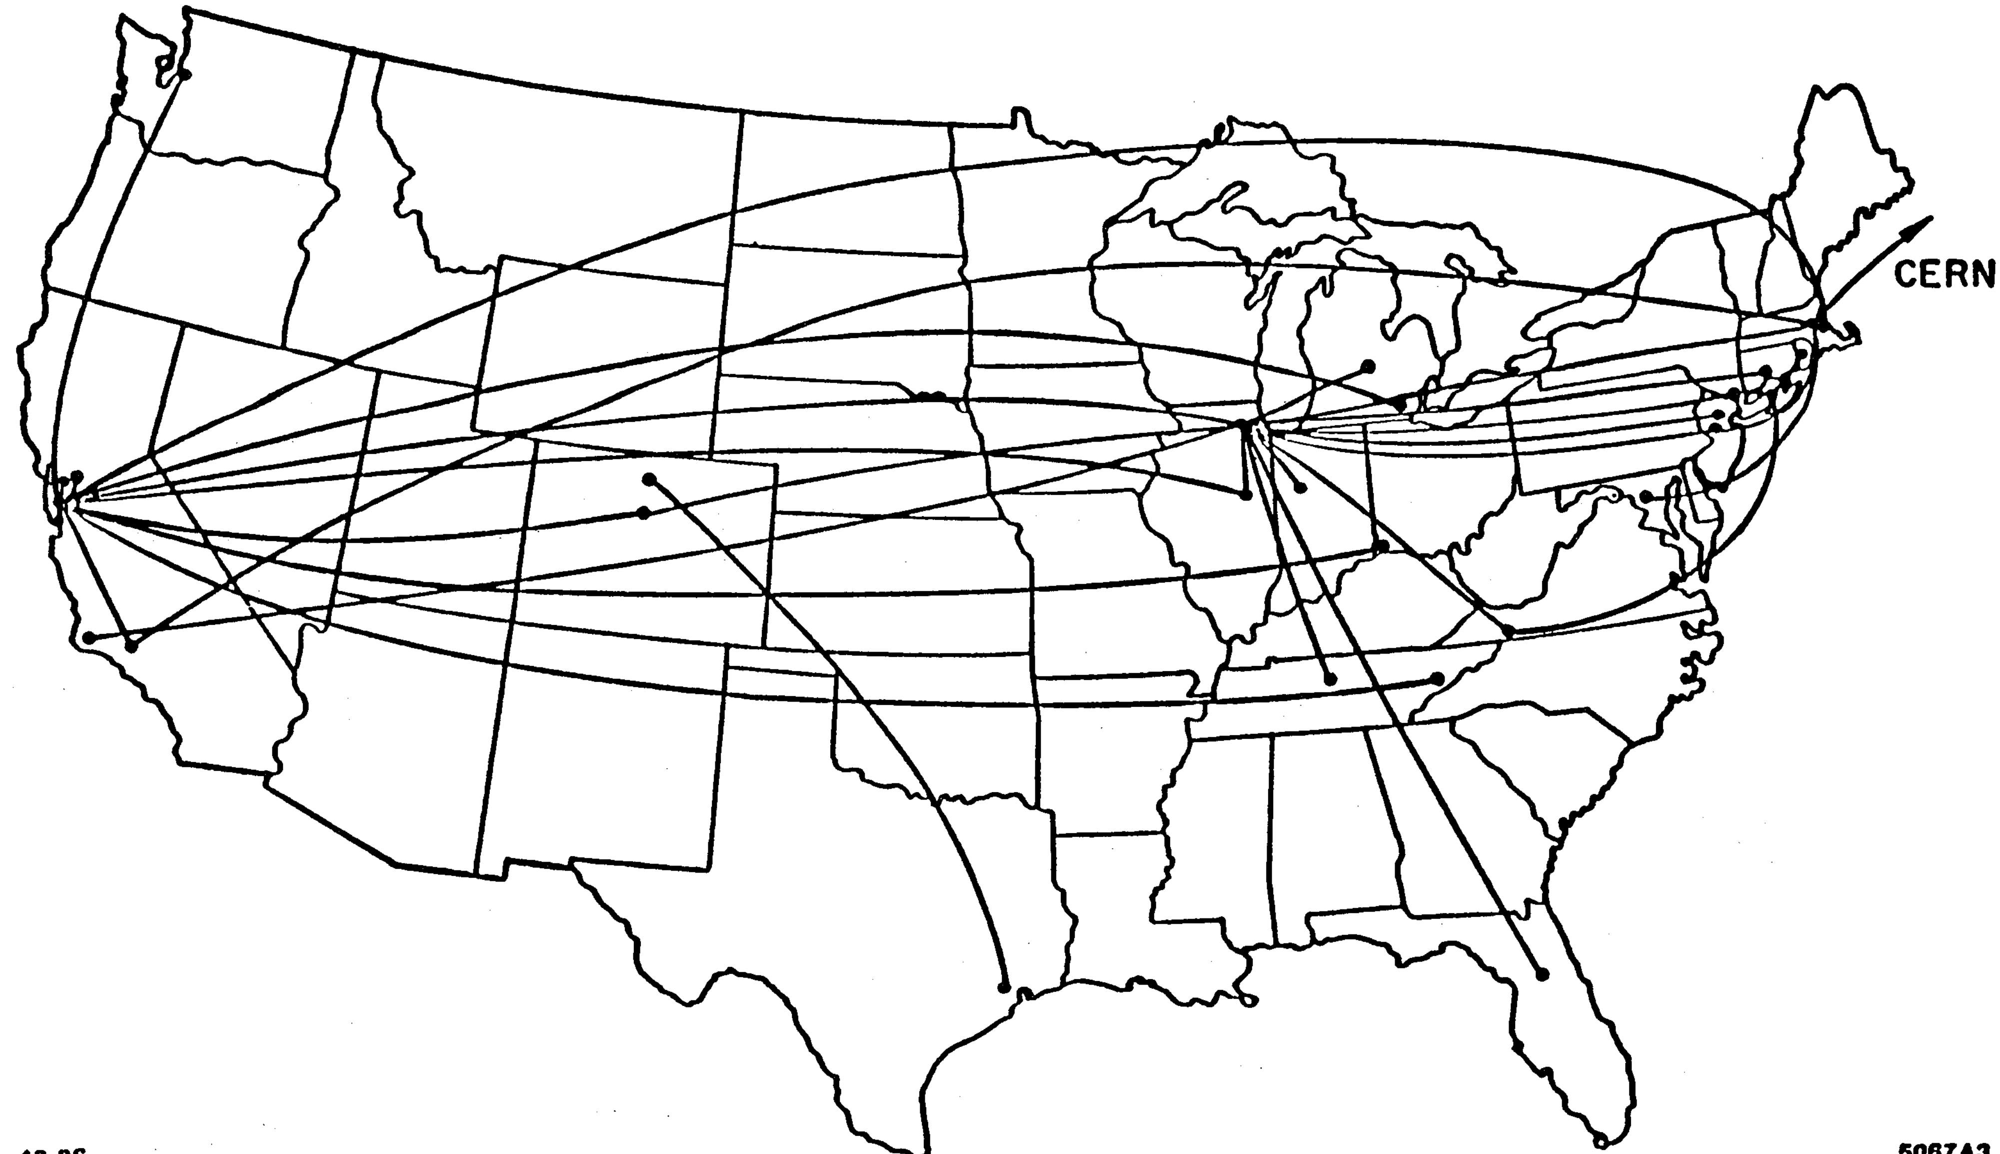

12-86

5067A3

FIG. VI.2.1

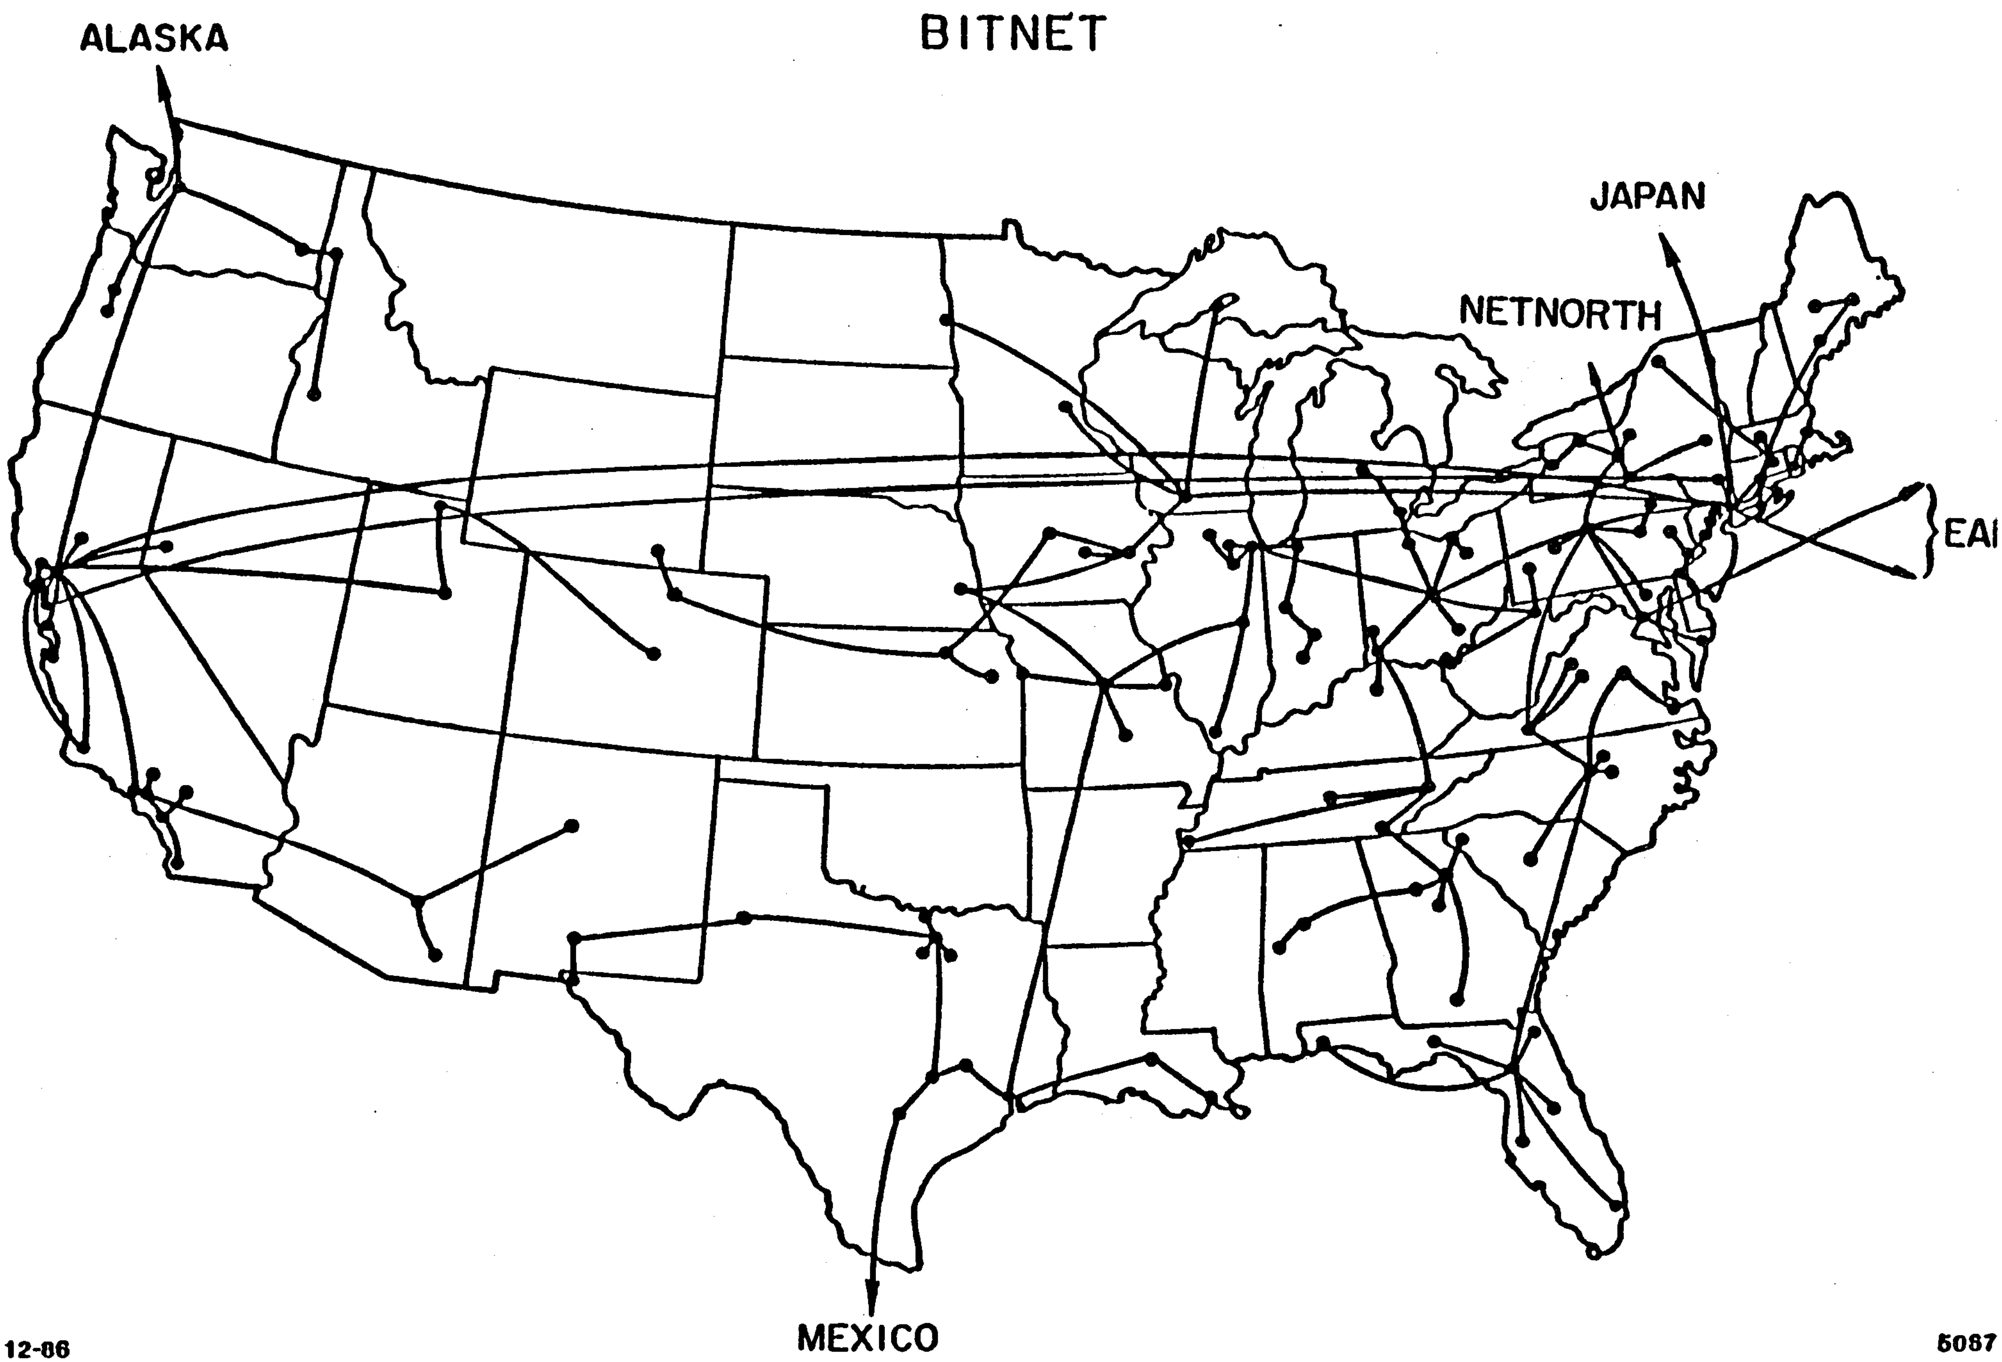

FIG. VI.2.2

DECNET Component of HEPNET (PHYSNET)

FIG. VI.2.3

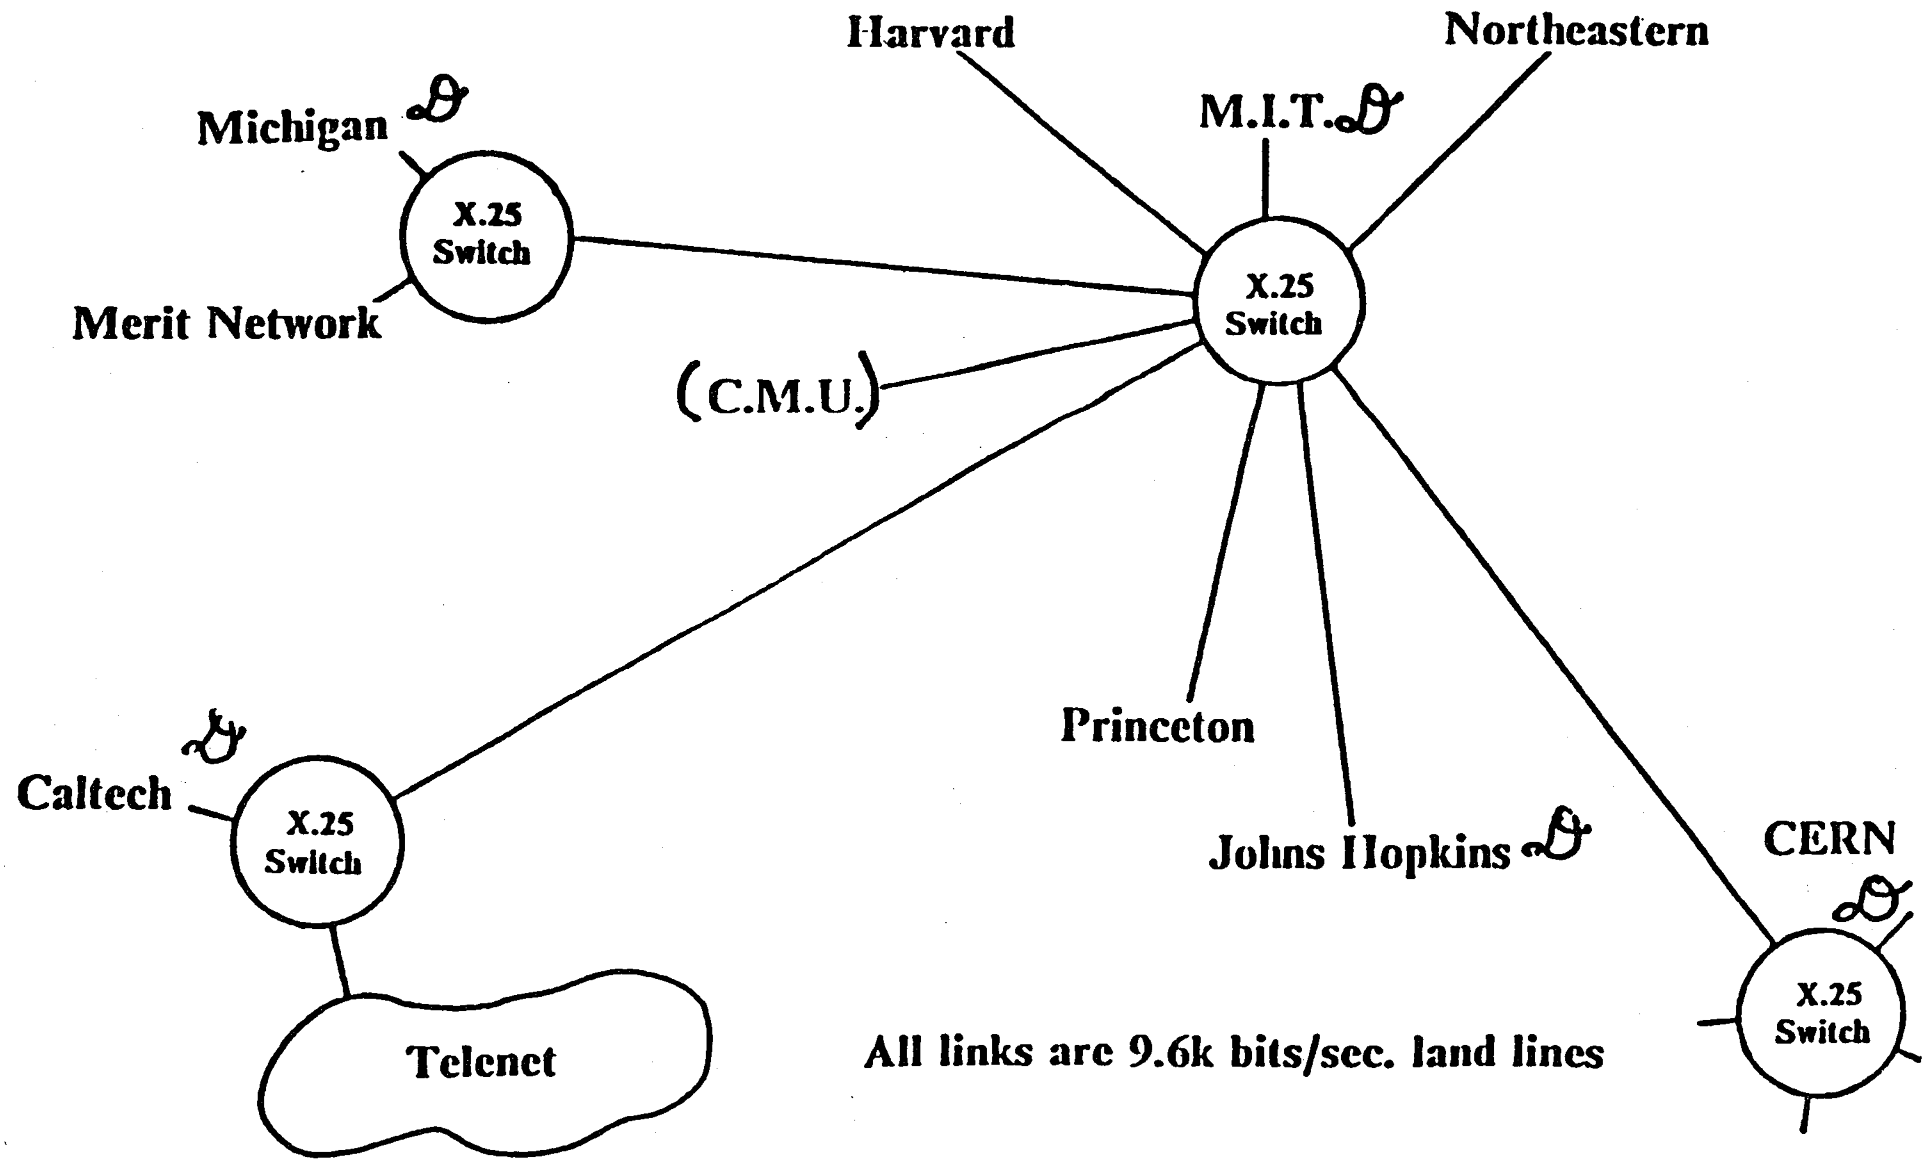

Harvard
Northeastern
Michigan
M.I.T.
X.25
Switch
X.25
Switch
Merit Network
(C.M.U.)
Princeton
Caltech
X.25
Switch
Johns Hopkins
CERN
X.25
Switch
Telenet
All links are 9.6k bits/sec. land lines

FIG. VI.2.4

TABLE VI.4.X.1

Leased Lines for High Energy Physics

Compliation by Paul Kunz <PFKEB0SLACVM>
25 November 1986

Amended by members of H.T.C.C.
Latest Revision 16 December 1986

To be included in this list the line must meet the following conditions:

- Line is used by HEP.
- Line goes off site or campus
- Lines are paid (at least in part) by HEP, Physics Dept, DoE, or NSF.

Notes:
- Many physical lines are used with multiple protocols, they are grouped together with a '/' and '\' near the line speed.
- Under 'type' of line: LL= Land Land, uW = MicroWave, Sat. = Satellite, and PPSN = Public Packet Switching Network.

From	To	Speed	type	Protocol	node	node	Use
ANL	U.Ill-Chi	9600	LL	RSCS	ANLOS	UICVM	BITnet
ANL	Purdue	4800	LL	DECnet	ANLHEP	PURDUE	DECnet
ANL	MFECC	56K	LL	MFEnet	?	?	MFEnet
ANL	U. Minn.	9600	LL	DECnet	ANLHEP	MINN	DECnet
ANL	Mich	9600	LL	DECnet	ANLHEP	MICH	DECnet
ANL	TymNet	9600	PPSN	X.25	ANNEX	terminals	Login
ANL	TymNet	9600	PPSN	X.25	IBMsw	terminals	Login
ANL	ArpaNet	56K	LL	TCP/IP	ANL-MCS	?	various

From	To	Speed	type	Protocol	node	node	Use
BNL	Yale	9600	LL	RSCS	BNL	YALEVM	BITnet
BNL	Brown	9600	LL	DECnet	BNLDOR	BRHEP1	DECnet
BNL	Boston U	9600	LL	DECnet	BNLRS2	BUPHYC	DECnet
BNL	Columbia	9600	LL	DECnet	BNLRS2	NEVIS	DECnet
BNL	Columbia	9600	LL	Statmux	Micom	Terminals	Logon(plan)
BNL	Cornell	9600	LL	DECnet	BNLDOR	LNS61	DECnet
BNL	John Hop.	9600	LL	DECnet	BNLRS1	JHUP	DECnet
BNL	NYU	9600	LL	DECnet	BNLDOR	NYUHEP	DECnet
BNL	U of Penn	9600	LL	DECnet	BNLDOR	UPENN1	DECnet
BNL	MFECC	9600	LL	MFEnet	BNLCL2	Princeton	MFEnet
BNL	SUNY-SB	9600	LL	Statmux	Gandalf	Terminals	Logon
BNL	SUNY-SB	/9600	LL	DECnet	BNLRS2	SBNUC1	DECnet
		\	LL	RSCS	SUNYSBNP	BNLDAG	BITnet

From	To	Speed	type	Protocol	node	node	Use
FNAL	Argonne	/48K	uW	RSCS	FNALVM	ANLOS	BITnet
		* 9600	uW	DECnet	CDFRT0	ANLHEP	DECnet
		* 9600	uW	DECnet	FNALR6	ANLPHY	DECnet
		\48K	uW	MFEnet	CCP/NAP	FNAL	MFEnet
FNAL	BNL	9600	LL	DECnet	FNALR0	BNLDOR	DECnet
FNAL	Columbia	/9600	LL	Statmux	Micom	Terminals	Logon
		\4800	LL	TCF	Cyber	printer	Printer
FNAL	U. Chicago	4800	LL	Statmux	Micom	Terminals	Logon
FNAL	U. Chicago	9600	LL	DECnet	FNALR4	UCHEP	DECnet
FNAL	U. Colo.	/1200	LL	Statmux	Micom	Terminals	Logon

		\4800	LL	TCF	Cyber	printer	Printer	
FNAL	U. Florida	/2400	LL	Statmux	Micom	Terminals	Logon	
		\9600		DECnet	FNALR4	UFHEP	DECnet	
FNAL	Harvard	9600	LL	DECnet	CDFRT0	HUHEPL	DECnet	
FNAL	U. Ill.UC	/1200	LL	Statmux	Micom	Terminals	Logon	
		* 2400	LL	DECnet	FNALR4	UIHEPA	DECnet	
		\4800	LL	TCF	Cyber	printer	Printer	
FNAL	U.Ill-Chi	9600	LL	RSCS	FNAL	UICVM	BITnet	
		9600	LL	DECnet	FNALR5	?????	DECnet	(1/87)
FNAL	U.Indiana	/4800	LL	DECnet	FNALR5	IND	DECnet	
		\9600		Statmux	Micom	Terminals	Logon	
FNAL	LBL	9600	Sat.	DECnet	FNALR4	LBLR1	DECnet	
FNAL	MIT	/2400	Sat.	Statmux	Micom	Terminals	Logon	
		\4800	Sat.	TCF	Cyber	printer	Printer	
FNAL	Mich St.	/4800	LL	Statmux	Micom	Terminals	Logon	
		* 2400	LL	DECnet	FNALR4	MUHEP	DECnet	
		\4800	LL	TCF	Cyber	printer	Printer	
FNAL	N.Western	9600	LL	DECnet	FNALR5	NUHEP	DECnet	
FNAL	Princeton	/1200	Sat.	Statmux	Micom	Terminals	Logon	
		\9600		DECnet	FNALR5	PUPHEP	DECnet	
FNAL	Rochester	9600	Sat.	DECnet	FNALR5	URHEP	DECnet	
FNAL	Rutgers	/2400	Sat.	Statmux	Micom	Terminals	Logon	
		* 9600		DECnet	FNALR4	RUTNPL	DECnet	
		\4800		TCF	Cyber	printer	Printer	
FNAL	UCSB	/1200	Sat.	Statmux	Micom	Terminals	Logon	
		\4800	Sat.	TCF	Cyber	printer	Printer	
FNAL	TymNet	9600	PPSN	X.25	Micom	Terminals	Logon	
FNAL	Vanderbilt	/4800	LL	Statmux	Micom	Terminals	Logon	
		\4800	LL	DECnet	FNALR4	VUHEP	DECnet	
FNAL	VPI	9600	LL	Statmux	Micom	Terminals	Logon	
FNAL	Wisconsin	4800	LL	DECnet	CDFRT0	PSLA	DECnet	
FNAL	YALE	/4800	LL	DECnet	FNALR4	YALPH2	DECnet	
		\2400		Statmux	Micom	Terminals	Logon	

From	To	Speed	type	Protocol	node	node	Use
SLAC	LBL	/9600	uW	RSCS	SLACVM	UCBCMSA	BITnet
		/		(Backup to BOC line)			
		* >56K	uW	DECnet	TPCS	LBLR1	DECnet
		\ 9600	uW	T1	Micom	Develcon	Logon
		\9600	uW	T1	TPCS?	Develcon	Logon
SLAC	UCB	9600	LL	RSCS	SLACVM	UCBCMSA	BITnet
SLAC	UCSB	/9600	LL	RSCS	SLACUCD	SBHEP	BITnet
		\9600	LL	DECnet	UCD	UCSB	DECnet
SLAC	Argonne	9600	Sat.	DECnet	HRS	ANLHEP	DECnet
SLAC	BostonU	9600	?	Statmux	Micom	terminals	Logon
SLAC	Caltech	/4800	LL	DECnet	SLD	CITHEX	DECnet
		/ 9600	LL	Statmux	Micom	Terminals	Logon
		\ 2400	LL	RSCS	SLACVM	CTECHPR1	Printer
		\9600	LL	RSCS	SLACVM	IMCTECH	Laser
SLAC	UC Davis	/9600	LL	Statmux	Micom	Terminals	Logon
		\2400	LL	RSCS	SLACVM	IMUCD	Laser
SLAC	TymNet	9600	PPSN	X.25	Micom		Logon
SLAC	MFECC	9600	LL	Statmux	GUSS	Micom	Logon
SLAC	U. Cinn.	9600	Sat.	Statmux	Micom	Terminals	Logon
SLAC	U. Colo	/9600	LL	Statmux	Micom	Terminals	Logon
		\2400	LL	RSCS	SLACVM	UCOLOPR1	Printer
SLAC	Fermilab	9600	Sat.	Statmux	Micom	Micom	Logon
SLAC	U. Ill.	9600	Sat.	Statmux	Micom	Terminals	Logon
SLAC	U. Mich.	/9600	Sat.	Statmux	Micom	Terminals	(planned)

		+ 9600	Sat.	RSCS	SLDMKII	UMICH	BITnet
		\9600	Sat.	DECnet	MKII	MICH	DECnet
SLAC	Northrdge	9600	LL	Statmux	Micom	Terminals	Logon
SLAC	SF State	/9600	LL	Statmux	Micom	Terminals	(planned)
		+ ????	LL	DECnet	??	?	DECnet
		\2400	LL	RSCS	SLACVM	IMSFSU	Laser(plan)
SLAC	UCSD	9600	LL	DECnet	UCD	SDPH1	DECnet
SLAC	UCSC	/9600	LL	Statmux	Micom	Terminals	Logon
		\9600	LL	RSCS	SLACVM	IMUCSC	Laser
SLAC	Stanford	9600	LL	RSCS	SLACVM	STANFORD	BITnet
SLAC	U. Tenn	/9600	LL	Statmux	Micom	Terminals	Logon
		\2400	LL	RSCS	SLACVM	IMUTENN	Laser
SLAC	U. Wash	9600	Sat.	Statmux	Micom	Terminals	Logon

From	To	Speed	type	Protocol	node	node	Use
Boston U	NorthEast.	9600	LL	RSCS	BOSTONU	NEUVMS	BITnet
CIT	Telenet	4800	PPSN	ClrBooks	CITHEX	various	L3NET
MIT	CERN	/16800	LL	ClrBooks	MITLNS	many	L3NET
		\9600	LL	DECnetX25	MITLNS	?	DECnet
MIT	Caltech	/9600	LL	ClrBooks	MITLNS	CITHEX	L3NET
		\9600	LL	DECnetX25	MITLNS	CITHEX	DECnet
MIT	Princeton	/9600	LL	ClrBooks	MITLNS	PRINHEP	L3NET
		\9600	LL	DECnetX25	MITLNS	?	DECnet
MIT	U Mich	9600	LL	ClrBooks	MITLNS	MICH	L3NET
	Merit	9600	LL	ClrBooks	MITLNS	various	L3NET
MIT	JHU	/9600	LL	ClrBooks	MITLNS	JHUPHEP	L3NET(plan)
		\9600	LL	DECnetX25	MITLNS	?	DECnet
MIT	Harvard	9600	LL	ClrBooks	MITLNS	HUHEPL	L3NET
MIT	NorthEast	9600	LL	ClrBooks	MITLNS	NORHEP	L3NET(plan)
MIT	CMU	9600	LL	ClrBooks	MITLNS	CMUHEP	L3NET(plan)
ColoSt	Houston	9600	??	Statmux	????	Terminals	Logon
Harvard	Brandeis	4800	LL	DECnet	HUHEPL	BRND	DECnet
Harvard	Telenet	1200	PPSN	X.25	?		L3NET
LBL	UCLA	4800	LL	DECnet	LBLR1	UCLA	DECnet
LBL	UCR	4800	LL	DECnet	LBLR1	UCR	DECnet
U. Minn	Soudan	9600	LL	DECnet	MINN	MINE	DECnet
UMich	Telenet	1200	PPSN	X.25	??		L3NET
U. Pa.	Princeton	9600	LL	RSCS	PENNDRLS	PUCC	BITnet
			(The computer is mostly owned by Physics Dept.)				
UCB	CUNY	9600	LL	RSCS	UCBCMSA	CUNYVM	BITnet
		(SLAC contributes a fraction of this line's cost)					
UCB	LBL	9600	LL	RSCS	UCBCMSA	LBLGATE	printer
PSL	Wisconsin	9600	LL	RSCS	WISCPSLB	WISCVM	BITnet
Tx AC	Houston	9600	??	Statmux	?	Terminals	Logon
MFECC	NRL	1200	??	MILNet	????	??	

DOE/ER-0234

COMPUTING FOR PARTICLE PHYSICS

Report of the HEPAP Subpanel on Computer Needs for the Next Decade

August 1985

U.S. Department of Energy
Office of Energy Research
Division of High Energy Physics
Washington, D.C. 20545

4. Analysis of Networking Requirements

4.1 INTRODUCTION

This chapter is a summary of the report of Working Group V which was charged with assessing the needs for low speed and high speed networking, and to evaluate the need for an integrated network for high energy physics.

Networks are already used extensively in high energy physics. Most of these have evolved on an ad-hoc basis as required for a specific experiment. Links are used either for direct terminal access to computer centers or for computer to computer services. In many cases, these networks have linked to include other institutions or other experiments.

An example of this pattern of growth is the network of DEC computers using DECNET. The PEP4 experiment (TPC) set up a microwave link between LBL and SLAC. All the VAX computers at SLAC were connected by cables in the PEP tunnel. LBL was linked to UCLA and Riverside by phone lines. The Two-Gamma experiment (PEP9) established links with Davis, Santa Barbara and San Diego. The HRS experiment leased a link to ANL, which was in turn connected to Michigan, Indiana, and Purdue. At Fermilab, CDF established a link to ANL, as well as links to Harvard, Brandeis, and Wisconsin. This provided a link from Fermilab, through ANL, to SLAC and LBL. Also at PEP, the DELCO experiment leased a line to Caltech.

Another example of the ad-hoc growth of networks is BITNET which links university computer centers, departmental computers, and research laboratories with 9600 baud leased lines. IBM is supporting, for a period of 4 years, a similar network in Europe, the European Academic Research Network (EARN). Two lines link

BITNET and EARN. Starting in 1980 with a line between the City University of New York and Yale University, the combined BITNET/EARN has grown to roughly 200 sites with 450 computers. Computers used by High Energy Physics groups make up about 15% of this network which includes SLAC, Fermilab, ANL, Cornell, BNL, CERN, and DESY, as well as universities in the United States, Canada (called NETNORTH), Europe, and Israel.

Networks will continue to be very important in high energy physics. They permit more effective use of manpower and computing resources. They facilitate communication between experimenters at their home institutions and those at the site of the experiment. With these links faculty and students can maintain close ties with an experiment and can help with data analysis and with diagnosing problems on a running experiment.

The experiments now under construction have immediate needs for networking. The networks will be used during the program development effort and later for the study of detector performance and analysis. Obvious examples include CDF, D0, LEP3, and SLD. The increased size and complexity of fixed target experiments will make networking more important for these smaller efforts as well as for non-accelerator based experiments.

4.2 Character of Network Use:

Networking includes a broad spectrum of services. These include the following:

- Direct terminal connection.
  Leased lines or direct lines within a local site are used to connect a terminal to a remote computer.
- Virtual terminal access.
  A local terminal on one computer is used to communicate with a remote computer over a computer-to-computer link.
- Mail.
  Computer to computer links are used to send mail messages to users on remote

computers.

- Phone.

  Computer links are used to exchange messages interactively with users who are logged in on a remote computer.

- File Transfer.

  Links are used to copy files to or from a remote computer.

- Remote Job Entry.

  A batch job is prepared on a local computer and submitted for execution on a remote computer.

- Remote Printing and Graphics.

  Files are transmitted from a remote computer for printing or display at a local device. It is also useful to transmit a file from a local computer for display or printing at another site. For these applications, the use of standard graphics and type-setting is extremely important.

- Distributed Databases and Libraries.

  Databases (e.g. calibration files) or program libraries are maintained on a single computer and are accessed from many other computers. By keeping only one version of the data or program, a large collaboration can more efficiently manage program development and analysis. With duplicate versions, variants can arise which must later be merged.

- Process to Process Communication.

  A program is executed on a remote computer on command from the local computer. For example, an experimenter at a home institution could start a process to look at new data from a detector and report any problems with the experiment. The networks are used to control an experiment from a remote site. It is unlikely that this will be standard operation for any experiment but it may be extremely useful in some circumstances. For example, an expert may be required to run special diagnostics from his home institution.

- Telefax.

  The network links are used to transmit facsimile copies. This could be a general case of remote printing or graphics.

- Video Conferencing.

  The links are used to transmit video images. These could include displays for a presentation.

4.3 Existing Networks:

A number of networks have already been established to satisfy the needs of experiments. These include direct terminal links and computer-to-computer service.

- Data switches:

  The digital switches at Fermilab, LBL, and SLAC have been linked by leased lines and microwave to provide terminal access between these three laboratories. These connections work at up to 9600 baud. Some universities also use these switches.

- DECNET:

  Network software from Digital Equipment Corporation, running on VAX or PDP-11 computers, has been used by numerous experiments to link their computers.

- BITNET:

  The BITNET software is available for IBM and DEC computers. A network of 9600 baud leased lines is used to link computers at many universities and laboratories. The network is used mainly for mail service and interactive messages, with some file transfer.

- X.25/COLOURED BOOKS:

  This network uses the X.25 Packet Switching protocol and connects computers of many manufacturers. The software system is a family of routines called coloured books.

- Other Networks:

  There are many other networks in the United States and elsewhere. None is used extensively by high energy physics. Most are available through gateway computers.

An example is the non-research part of ARPANET, now called MILNET.

None of these networks provides all the features desired for high energy physics. In Table 4.1, the features of the available networks are summarized.

4.4 CURRENT AND FUTURE USE:

In responding to the questionaire from the Subpanel, (see Chapter 2) most high energy physics groups indicated that they already use some network link to access laboratory resources or to communicate with their collaborators. A list of high energy physics sites which are on networks at this time is given in Table 4.2. Within the United States, the network links are used for terminal access to central computers and for computer-to-computer communication.

The diversity of network use reflects a variety of approaches to computing. These approaches can be categorized in general as centralized or distributed computing. Many groups use only a central computer for the analysis of an experiment. Direct lines on site or leased lines are used to provide terminal access to the computer or to the central terminal switch. The program library can be managed at the central site. In many other experiments, the computing is done at widely distributed sites. The program library is distributed to all computers on magnetic tape or by file transfer. Most groups who have adopted this style have established computer-to-computer links such as DECNET. Any future networks must support both approaches to computing.

The existing links are shown on maps in Figures 4.1, 4.2, and 4.3. There are a number of institutions with links to at least two of the laboratories. A more coherent approach would remove the necessity for these extra lines. In addition, a coherent network would provide communication with groups outside the individual collaborations and would encourage smaller university groups to establish a link to one of the laboratory sites.

Within the next four years new large colliding beam detectors at SLAC, Fermilab and Cornell will become operational. At the same time a new generation of detectors and experiments will be forthcoming at the TEV II fixed target program at Fermilab. All of these efforts need extensive networking now to improve communication during the design and program development phases. When these experiments begin to take data, the networks will be used for distribution of sample data and for the communication of new results. The speed requirements will increase as the networks are more heavily used.

In addition to networking to support the experimental program, there is an increasing demand for access to the MFE Computing Center to support Theoretical and Accelerator Physics. In particular, the Central Design Group for the SSC uses a 9600 baud leased line to LBL. This line may be upgraded in the future as the SSC use of the CRAY increases.

The speed requirements for networks are estimated by considering some examples for services to be provided. Table 4.3 illustrates some typical transfer that might be expected. In the short term, it seems that 9600 baud is sufficient for all direct links. Trunk lines between laboratories must carry data for many of the direct links and will require 56k baud initially. In the next three to five years, the volume of data carried by the trunks is expected to increase by a factor of three to five.

The development of a High Energy Physics Network (HEPNET) is motivated by the need to provide network services to a large fraction of the high energy physics community and to reduce the total cost of those services. Cost savings can be realized by eliminating links to two of the laboratories. For example, there are currently five lines leased by high energy physics groups to connect the West Coast to the Midwest. New programs will require at least three more in the next year. All of these groups are also connected to SLAC or LBL so that a link to Fermilab from SLAC/LBL would satisfy the requirements of all of them. The cost of a 56k baud link is approximately four times that of individual 9600 baud lines so that the total cost could be reduced by a factor of two. In addition, groups whose requirements do not justify two links will have access to all sites for the cost of one connection. A similar situation for East Coast groups justifys a dedicated trunk to the Midwest.

4.5 MODEL FOR HEPNET:

The plan for a dedicated High Energy Physics Network (HEPNET) is based on two important constraints.

The functionality and quality of all existing links must be maintained. The terminal lines and other links are necessary for the physics program. Any new links must work as well.

HEPNET must provide the following services to all users:

- Direct terminal access to all computer centers.
- All DECNET functions
- All BITNET functions
- X.25 (especially for communication to Europe and Japan)

The network is based on high speed links between the high energy physics laboratories in the United States. Links from LBL/SLAC to Fermilab, Fermilab to BNL, and BNL to LBL/SLAC at 56 kbaud will provide sufficient bandwidth and redundancy. As needed, these links should be upgraded.

The trunk lines can be either leased land lines or satellite links. The choice between these two techniques will be based on economic factors and on the need to provide good interactive response for terminals on remote computers. The diverse requirements may be best satisfied by a combination of the two techniques.

University groups will use feeders to communicate with the trunk. Each group will choose the site of highest traffic, not necessarily the closest node. The costs of leased lines are rather insensitive to distance so that the choice will not have a significate impact on the total cost.

The cost of the network can be estimated rather reliably. Most of the cost is in the charges for leased lines. The cost of three satellite trunks is approximately $120,000 per year. Feeders for 60 universities will cost about $350,000 per year. The exact cost of the hardware at each network node will depend on the protocol adopted but will be significantly less than the annual line costs. Land lines are roughly twice the cost of satellites.

4.6 FUTURE NETWORKS:

A number of new projects are underway to provide additional networking within the scientific community. While not addressed to the problems of high energy physics, these approaches may be useful in the long term.

SCIENCENET is a network designed to provide access to supercomputer sites established by the National Science Foundation. There are no explicit plans to establish links to the high energy physics laboratories.

The DoE will begin using OPMODEL, a dedicated satellite for high speed communication. The satellite could, in the long term, provide the trunk lines suggested above at a significant saving to the high energy physics community.

We don't know whether either or both of these projects may become cost effective alternatives to the leased lines for a dedicated HEPNET. In the short term, however, they will not be available and the development of a High Energy Physics Network should proceed independently. The management of HEPNET must maintain close communication with these efforts to provide gateways or share their resources wherever possible.

4.7 CONCLUSIONS:

- Networks are extremely important to the high energy physics program.
- Some experiments have used networks extensively and these can be used as a foundation for a more complete High Energy Physics Network.
- Significant cost savings can be realised by a more coordinated approach to networking. The network can be made much more useful to all users at the same cost. In particular, it will be important to establish links to groups outside individual experiments.

DOE/ER-0341 ESNET June 1987

UNITED STATES
DEPARTMENT OF ENERGY
WASHINGTON, D.C. 20545

ER-7