DOE F 1325.8
(4-84)

United States Government | Department of Energy

memorandum

DATE: OCT 07 1986

REPLY TO ATTN OF: ER-2

SUBJECT: ER Supercomputer Strategy

TO: James F. Decker, ER-2

I concur with the recommendations which you present in your memorandum dated September 18, 1986. The acquisition of a Class VII Supercomputer in FY 1988 to support ER scientists has been given high programmatic priority. In view of the economic analysis presented, I agree that this supercomputer system should be installed at the National Magnetic Fusion Energy Computer Center (NMFECC) and that the NMFECC scope should be expanded to be a permanent ER-wide supercomputer center. The establishment of additional ER supercomputer centers should be considered in the future as demand for supercomputer resources and budget factors indicate.

I also agree that the Scientific Computing Staff should move forward to implement the Energy Sciences Network. Funding for the supercomputer access program will continue to be provided through the budget and reporting codes in the Offices of Fusion Energy and Basic Energy Sciences which currently support this activity. However, the Scientific Computing Staff will assume budgetary responsibility and control of the activities under these budgets.

Please continue to keep me informed with regard to the progress of these important projects.

Alvin W. Trivelpiece
Director, Office of
Energy Research

cc: R. Young, ER-60
I. Adler, ER-61
D. Mayhew, ER-63
W. Wallenmeyer, ER-20
J. Cavallini, ER-7
D. Stevens, ER-10
C. DeLisi, ER-70

SEP 2 9 1986

ER-2

ER Supercomputer Strategy

A. W. Trivelpiece, ER-1

Recently, the Scientific Computing Staff prepared the attached analysis of the ER supercomputing requirement and long term strategy for addressing this requirement. This analysis, although strongly emphasizing the need for additional supercomputer capacity and capability as well as the desirability of operating more than one supercomputer center to provide backup capability, alternative architecture support and broader software support, recognizes the fiscal issues related to budget deficit reduction as the paramount issues during the next several years.

The magnitude of the demand for supercomputing to tackle difficult ER computational problems and the significant accomplishments that have already resulted from the creation of the ER supercomputer access program firmly establish the critical importance of acquiring a more advanced supercomputer system for the ER user community in FY 1988. However, the fiscal issues dictate that this advanced system be acquired and operated as cost effectively as possible in the near term. The analysis also shows that the most cost effective way to provide support for such a system is to continue to use the NMFECC, since the NMFECC will need only incremental base support increases to install and to operate an additional system.

Therefore, I recommend that the NMFECC be designated as the permanent ER supercomputer facility to support any new system acquired in FY 1988 while recognizing the desirability of operating an additional supercomputing facility when budgets allow.

I have also attached an analysis of the data collected as a result of the ER data communications survey. This analysis verifies that an integrated approach to ER data communications can provide significant benefits. As a result of this and other initiatives, the Scientific Computing Staff proposed an implementation plan for evolving the MFENET into an Energy Sciences Network (ESNET) as a general purpose scientific network for the ER community. This proposal has received positive response from the ER programs and from other federal research programs. I recommend that we proceed vigorously with the implementation of the ESNET.

If you concur with these recommendations, please sign the attached memo and I will initiate the necessary actions to implement them.

Signed by
James F. Decker
James F. Decker
Deputy Director
Office of Energy Research

Attachments

ER-7:JCavallini:bc:353-5800:9/18/86

bcc: ER-1(info);ER-2;ER-60;ER-622

CONCURRENCE

RTG. SYMB	INITIALS/SI	DATE
ER-7	Cavall	9/18/
ER-60	Young	9/ /8
ER-2	Decker	9/19/8

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 proposed 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.

Energy Research
Supercomputer Requirement
and
Long Term Strategy

Background

During FY 1984, the Office of Energy Research responded to several reports, both internal and external, which found that the availability of modern supercomputer resources to the U.S. scientific research community fell far short of the amount needed. Subsequently, the Energy Sciences Advanced Computation activity was initiated to establish an ER-wide supercomputer access program. The strategy adopted at that time was to take advantage of existing facilities within the Fusion Energy program. This was done by expanding the role of the National Magnetic Fusion Energy Computer Center (NMFECC) and its nationwide data communications network to service researchers in the other ER programs, i.e., High Energy and Nuclear Physics, Basic Energy Sciences and Health and Environmental Research. Initially five percent of the resources at the NMFECC were made available to these other ER programs. During FY 1985, a Cray X-MP/22 computer system was acquired on an interim basis to better address this need until the ER supercomputing requirement could be properly sized, validated and documented and until a permanent site could be decided upon to support this requirement.

Requirement

Research at the forefront of contemporary and future science and technology as done within the ER programs will demand adequate access to supercomputer power. This fact has been verified in several ways. First, the report, 'The Role of Supercomputers in Energy Research Programs,' published in

February 1985, identified and documented many problems throughout the ER program areas which require supercomputing resources. This report, which was compiled by researchers from all ER programs, also identified many problems for which more advanced supercomputers will be necessary. Second, this need for supercomputing resources has been verified by the high demand for and especially the high level of usage of these resources by researchers in all ER programs. The high demand is exhibited through proposals for resources submitted to ER headquarters from the ER research community. The following table summarizes the large growth rate of this demand versus availability for supercomputing resources for the non-fusion ER programs:

	FY 84	FY 85	FY 86
Hours Requested	10,400	24,080	45,060
Hours Available	800	12,000	24,000

The high demand for resources by itself would not substantiate this requirement, however, ER researchers have substantiated this requirement by completely utilizing all of the resources that have been made available to them.

The last way that this requirement has been verified is through the benefits to the ER programs that supercomputing has made in a very short time. Significant accomplishments have been made in all program areas. In High Energy Physics, researchers have been able to realistically model Klystrons for the first time. In Materials Science, researchers have been

investigating new superconducting materials. In Applied Math, researchers have developed new algorithms for supercomputers to solve ER computational problems, e.g., an algorithm to address the track finding problem in HEP and which operates orders of magnitude faster than scalar algorithms. In Health and Environmental Research, supercomputers are being used to model faunal buildup and to predict RNA secondary structures.

To summarize, this supercomputer requirement has been well validated through good usage and through documented accomplishments. Furthermore, this requirement is growing at a very high annual percentage. At current growth rates, this ER access program will be able to address less than ten percent of the ER supercomputing requirement within the next three years.

More importantly, many difficult problems cannot even be undertaken using current supercomputer systems. The amount of complexity incorporated in computational models is often scaled down to meet the existing systems capabilities. Additional supercomputer capability is needed to incorporate more dimensions, finer resolutions, more physics and longer time scales into ER computations. The real world exists in three space dimensions plus time, however, current systems for the most part can only treat ER computations of two dimensions plus time. In accelerator physics, magnet quality and alignment tolerance computations require much longer time scales to correctly simulate misalignment and fabrication errors and to estimate their influence on beam quality at the interaction point. In materials science, a critical input for solidification computations is the interparticle

interaction potential and faithful simulation of such potentials depends upon several thermodynamic state variables or physical effects, i.e., density, temperature, pressure. The coupled complexities of these physical effects of size and interaction potential computations are prohibitive on current systems. Therefore, the most advanced and capable systems available will be needed to support the ER supercomputer requirement. In particular, a Class VII system will be needed in FY 1988.

Supercomputer Facilities

Currently, ER operates only one Center which provides supercomputer resources, the NMFECC. ER researchers also are allocated some resources via a cooperative agreement with the Florida State University, a Congressionally initiated university project. For the past two fiscal years, ER has requested funds to provide permanent facilities to address this requirement. Because of budget constraints, these requests were denied. Nevertheless, the requirement is large enough to justify individual supercomputers within both the HENP and BES programs and there is a great deal of interest within the ER Laboratories for operating these facilities for this access program.

The basic issue in this regard is: Whether to redefine the mission of the NMFECC to be the overall ER supercomputer center permanently or not, i.e., to create an additional center. Some of the pros and cons are as follows:

Pros: 1. NMFECC has already incorporated all of the new ER users into the NMFECC user community.

2. The NMFECC data communications network has been expanded to provide access at all major ER locations.
3. The NMFECC staff has been modestly expanded to support this new requirement.
4. Fusion and non-fusion users will be able to swap resources across supercomputer systems and thereby promote software and resource sharing.
5. The cost for operating a single center (at present, the NMFECC) for this purpose and at this level of capacity has been proven to be less expensive due to shared facilities and manpower.

Cons:
1. NMFECC may need to expand its facilities to permanently house the additional staff.
2. There are potential cost recovery issues/problems in funding the NMFECC through two separate budget activities.
3. An additional supercomputer center would serve as a backup recovery center in case of disaster, such as earthquake.
4. A second center could provide access to alternative supercomputer architectures which would otherwise extend the NMFECC beyond its resources.

5. Healthy inter-laboratory competition for supercomputing expertise.
6. An additional center would broaden the software support and knowledge base for supercomputer systems and their use.

The most important considerations with regard to providing new supercomputer capability or additional support facilities for supercomputers revolve around the tightening fiscal constraints in the coming years and the political issues surrounding the FSU Institute. Additional funds required to create and support a second supercomputer center will be very difficult to obtain. At this point, since the NMFECC has already expanded staff, data communications, and file management support to handle the additional non-fusion users for FY 85, FY 86 and FY 87, NMFECC may only need to add an extension to the existing building to house the additional staff who are now in temporary facilities and to budget any needed increases for inflation. Based on these important issues then, the NMFECC appears to be the best candidate for assignment of this additional responsibility permanently. At present, the responsibility for operational management of the NMFECC at DOE Headquarters is assigned to the Scientific Computing Staff (SCS). The responsibility for resource allocation is split between the Scientific Computing Staff, which coordinates all non-fusion requests, and the Office of Fusion Energy (OFE) which handles all of the OFE allocations. Budget responsibility is also split between SCS and OFE for the NMFECC and this situation has resulted in several problems over the past two years. Since

budget formulation and priorities differ across ER programs, deficiencies have had to be corrected for the NMFECC operation in each budget cycle. Additionally, because NMFECC is supported by two funding activities, the NMFECC may not meet with OMB ADP cost recovery guidelines and may be required to institute a billing or charge back mechanism. The recently issued OMB circular A-130 appears to relax this cost recovery issue, but budget formulation problems may persist for the NMFECC because of the differing program priorities. Therefore, it is recommended that the management responsibilities for the NMFECC be consolidated within the SCS while recognizing past OFE contributions to the NMFECC. With regard to budgetary responsibility, it is critical to have close coordination between OFE and SCS to ensure stable operation at the NMFECC. Therefore, it is recommended that OFE develop its budget for this activity at a level consistent with its priorities but in close consultation with concurrence by the SCS. This level of formal consultation will improve the fiscal and budgetary management of the NMFECC while allowing OFE to determine the level of computing support it requires.

Cost Analysis

The FY 1987 budget for the NMFECC is $25.5 million, $16.0 million in OFE and $9.5 million in ESAC. An examination of the annual budgets of other supercomputer centers within the DOE shows that the cost of operating each state-of-the-art facility averages between $15 and $20 million annually over and above lease/mortgage expenses for the supercomputer system. The current Argonne National Laboratory institutional plan proposes to establish an additional ER supercomputer center and projects a budget of $18.5 million

for this center. Argonne also reduces the scope of the center by assuming that NMFECC or the vendor will be the primary software support and that NMFECC will continue to operate all nationwide data communications for the access program. It should be noted that the NMFECC currently supports two supercomputer leases at about $7.5 million and is a very cost effective operation.

For FY 1988, the additional funds needed to support a Class VII computer system at the NMFECC will be approximately $5 million, whereas the initial funding requested by Argonne for FY 1988 is $10.0 million. For FY 1989, the increase to support a full year lease-to-ownership at NMFECC is an additional $1.5 million but at Argonne the increase is $6.0 million. Hence, the savings by continuing to support this access program through the NMFECC are $5.0, $8.5 and $10.0 million for FY 1988, FY 1989 and FY 1990, respectively.

Recommendation and Long Term Strategy

The ER supercomputer access program has proven to be a vital tool for tackling many of the forefront problems facing ER scientists. The most advanced computer systems available in the marketplace should be made available for use by the ER research community through this program.

The magnitude of the demand for supercomputing for ER computational problems represents a requirement for multiple supercomputer systems and the diversity of ER problems represents the basis for utilizing differing supercomputer architectures. The need for a backup/recovery site for the NMFECC, the need to further our knowledge base for supercomputing, the benefits of broadening the supercomputer software base, and the benefits of interlaboratory competition and information exchange adequately justify the establishment of an additional supercomputer center within ER.

Nevertheless, the fiscal and political issues regarding budget deficit reduction and Congressionally initiated university projects will be the paramount issues during the next several years. Hence, the deciding point here should be cost effectiveness and, as shown in the previous section, continuing to use the NMFECC to provide the additional supercomputing capabilities which are needed by the ER community is the most cost effective approach.

Hence, it is recommended that the NMFECC be designated as the permanent ER supercomputer center and that the Class VII computer system be approved and be targetted for installation at the NMFECC in FY 1988. Because of the many benefits that can be obtained from operating an additional supercomputer center, it is recommended that this alternative be reconsidered in the FY 1990 or FY 1991 budget cycle.

This strategy, therefore, assumes the very conservative approach of addressing less than twenty percent of the requests for supercomputer resources that are projected through FY 1990. The strategy also assumes a consolidation of the management of ER supercomputing budgets to avoid potential problems during budget formulation and justification. If this access program results in the emergence of a pressing requirement within a specific ER program, e.g., SSC data analysis, then the strategy should be to finance a dedicated supercomputer resource through existing program funds and to manage the activity through the Scientific Computing Staff.

Lawrence Berkeley Laboratory
1 Cyclotron Road Berkeley, California 94720
(415) 486-4000 • FTS 451-4000

May 1, 1986
CDO-86-156

Mr. John Cavallini
U. S. Department of Energy
Office of Energy Research
Energy Sciences Advanced Computation
Washington, D. C. 20545

Dear John:

Attached is the preliminary report regarding the network summary. As we discussed, it is difficult to draw quantitative conclusions from the summary at this time.

I believe that an annual "Workshop on ER Computing Needs" is an excellent idea! It would provide an appropriate method of developing directions not only for networking needs but for whatever current questions need to be answered. I feel that this is needed now and, to that end, offer LBL as a host for the first workshop.

I look forward to working with you on these issues.

Sincerely,

Leroy T. Kerth
Associate Director
for Computing

LTK:da

D. F. Stevens, LBL

A Preliminary Summary of Responses to the OER Questionaire on External Data Communications

Executive Summary

The responses to the survey do not give enough information to support quantitative conclusions at this point. (See Appendix B for a discussion of the difficulties.) More information will be collected. The qualitative conclusions that may be drawn from the responses are:

1) Nearly all programs make extensive use of more than a few different data networks and communicate with a large number of other sites. The principal networks are MFENet, HEPNET (DECNET), BITNET, and ARPA/MILNET.

2) The OER programs require the high-level functionality provided by these different networks.

3) An integrated approach at the link level can lead to significant savings when implemented as a common carrier of a number of different higher-level protocols. ES inter-networking, interagency networking, and supercomputer access will be best served by a common network-level protocol.

These conclusions are discussed more fully below, and tabular summaries of the responses are provided in Appendix A.

1. **Nearly all programs make extensive use of more than a few different data networks and communicate with a large number of other sites. The principal networks are MFENet, HEPNET (DECNET), BITNET, and ARPA/MILNET.**

   ER scientists are desperate for convenient access to computers, colleagues, and data located at other sites. Of the twelve responses received, only one (ORAU) indicated negligible need for external data communications. The other respondents indicated that an aggregate of thousands of messages and files are transmitted and hundreds of remote terminal hours are spent daily. The desperation is evident in the fact that so many channels -- at least fourteen networks were reported -- are used, with a majority of sites reporting use of more than four services. More than four services means more than four different sets of protocols to accomodate. Scientists are not eager to complicate their working lives in this fashion, but communication is clearly more important than simplicity.

   All of the multi-program sites responding indicated that several different ER programs were significant users of external data communications.

   The exact number of sites that exchange ER-sponsored or ER-related information is not obtainable from the responses, but it certainly exceeds 100 locations in the US, and includes in addition several in Europe and Japan. CERN, in particular, is a major destination for external data communications for several DOE Laboratories. The US locations include all DOE National Laboratories, all major universities, and a number of contractors.

2) **The OER programs require the high-level functionality provided by these different networks.**

3) **An integrated approach at the link level can lead to significant savings when implemented as a common carrier of a number of different higher-level protocols. ES inter-networking, interagency networking, and supercomputer access will be best served by a common network-level protocol.**

   Three fundamentally different kinds of external data communications services are mentioned in the responses: remote interactive access to computers or information services at other sites, the transfer of large blocks of data from one site to another, and message-oriented services.

   i) Remote interactive access: There are two general situations in which a scientist seeks remote computing access to another site: To achieve access to a computer with capability beyond that which is available locally, or to achieve access to data that, because of currency or quantity, cannot

be transmitted from its place of residence to that of the user. In the former case, remote access replaces hardware acquisition; in the latter, it replaces travel.

ii) File transfer: This is in some sense a generalization of a message service, where the "message" is a program or data file, and where the destination is often a program or system utility (PRINT, for example) rather than a person. The user interface for file transfer is usually very different from the message interface on the same network, so that the user wishing to use both must learn two distinct sets of commands.

iii) Message-oriented services:

a) Person-to-person messages (E-mail): This type of service sees heavy use on all networks. It is the principal reason for the widespread implementation of BITNET. Many members of the extended ER community are accessible (for electronic communication) through no other means.

b) Bulletin Board services: We believe that bulletin board traffic is significantly under-reported in the survey responses. Users tend to see bulletin boards as a *local* service, and fail to realize that essentially all bulletin board messages either originate off-site or have at least one off-site addressee. Bulletin boards exist at all levels of specificity and practicality. The most common usage is to transmit system news about the host site, but there are also discipline-specific bulletin boards, and an increasing number of useful commercial bulletin boards (e.g., Autocad; Byte Magazine bulletin board).

c) Electronic conferences: These are generalizations of both bulletin board services and E-mail. They are more interactive than most bulletin boards and have longer memories than either bulletin boards or E-mail services. They contain utilities to assist users to look up old submissions, to vote, to contribute, to engage in private dialogue with other conferees, etc.

Appendix A: Tabular summaries

Table 1: Summary by Site

Site	ER %	Cost(K$)	ARPA	BITNET	HEPNET	MFENet	TYMNET	USENet	Others
Ames	100	2.5		x		x			
ANL	90	233.2	x	x	x	x	x	x	Argonne File Transfer Net
BNL	59	75.0	x	x	x	x		x	WETNet to Natl Weather Svc
FNAL	100	135.5		x	x	x	x		
ITRI	100	*Unspec.*					x		Telenet, UNINet
LBL	90	256.3	x	x	x	x	x		*Unspecified*
LNS	100(?)	0.6		x		x			CHAOSNET, LEP3NET
MFE	100	1580.0	x			x	x		
ORAU	0								*Negligible Traffic*
ORNL	31	513.4		x		x	x	x	DDN, BETANET
PNL	?	*Unspec.*				x			
SLAC	100	113.0		x	x	x	x		

Table 2: Summary by Network

(Only those networks mentioned by three or more sites are included; *F* = File Transfer Service, *M* = Message/Mail Service, and *T* = Terminal Access Service.)

Network	Site	Service	Connection	Principal Destinations
ARPA/MILNET	ANL	F, M, T	IMP, Dial-up	LLNL, LANL, LBL, BNL
	BNL	F, M, T	IMP	*Unspecified*
	LBL	F, M, T	IMP	Washington DC, ORNL
	LNS	F, M, T	via CHAOSNET	*Unspecified*
BITNET	Ames	F, M	Coax to IU	*Unspecified*
	ANL	F, M	Leased Line	DOE Labs, Major Universities
	BNL	F, M	*Unspecified*	CERN, DESY, Universities
	FNAL	F, M	Microwave to ANL	CERN, Member Universities
	LBL	F, M	*Unspecified*	CERN, SLAC
	LNS	F, M	Leased Line	*Unspecified*
	ORNL	F, M	via DDN	Universities and Laboratories
	SLAC	F, M	Leased Line	*Unspecified*

Table 2: Summary by Network (continued)

HEPNET/PHYSNet	ANL	F, M, T	Leased Line	SLAC, FNAL, UMich, Purdue, UInd
	BNL	*Unspec.*	*Unspecified*	FNAL, SLAC, Universities
	FNAL	F, M, T	Leased Line, Microwave to ANL	BNL, LBL, SLAC, Other Laboratories
	LBL	F, M, T	Leased Line, Microwave to SLAC	BNL, FNAL, SLAC
	LNS	F, M, T	via LEP3NET	*Unspecified*
	SLAC	F, M, T	Leased Line, Microwave to LBL	FNAL, LBL, Universities
MFENet	Ames	F, M, T	NAP	MFE
	ANL	F, M, T	CCP	MFE
	BNL	F, M, T	NAP	MFE
	FNAL	F, M, T	Microwave to ANL	MFE
	LBL	F, M, T	NAP/CCP	MFE
	LNS	*Unspec.*	*Unspecified*	MFE
	NFE	F, M, T	USC/CCP/NAP	All ER sites
	ORNL	F, M, T	Satellite Link	MFE, CBAF, Auburn, MDD, UIll, UWisc, FSU
	PNL	*Unspec.*	*Unspecified*	MFE
TYMNET	ANL	T	Dial-up	*Unspecified*
	FNAL	T	Dial-up	Member Universities, Collaborators
	ITRI	*Unspec.*	*Unspecified*	*Unspecified*
	LBL	*Unspec.*	Mini-Engine	New Jersey, Arizona, New York, DC
	ORNL	T	Net Host @ ORNL	RECON users
	SLAC	T	*Unspecified*	*Unspecified*
UUCPNet/USENet	ANL	F, M	Dial-up	Major Universities
	BNL	F, M	*Unspecified*	*Unspecified*
	ORNL	F, M	via DDN	*Unspecified*

Appendix B

The source data is insufficient to support quantitative conclusions for the following reasons:

1: The sample is too small. There are other sites yet to be polled.

2: FNAL, LBL, and SLAC responded in different terms for different types of services (*hours/day* of terminal access, *messages/month* for mail, *files/day* for file transfer, and *commands/day* for remote command execution); all other respondents giving traffic information responded in terms of *hours* for all types of service.

3: The use of *hours* as a unit does not allow one to distinguish file and message access from terminal access, or to distinguish between links of different bandwidths.

4: There may be a significant amount of duplication in the responses. This could happen not only as a result of the reporting of the same link by the sites at either end, but also as a result of reporting both a network and its physical medium. HEPNET, BITNET, and TYMNET, in particular, are frequently accessed through several different types of physical link (leased line, dial-up, microwave); it is possible in these cases that the same traffic is counted twice.

5: It is not clear whether MFENet entries refer to MFE computation allocation units, machine hours, or data traffic hours. ANL and LBL, for instance, have similar allocations of MFE time, yet their responses yielded extremely different values for MFENet traffic.

6: In some cases the data submitted did not include all the ER programs at a given site.

If quantitative information is desired it will be necessary to conduct a more complete survey, and to provide more detailed instructions on the kinds of access and traffic that are of interest, and the reporting units that should be used for each type of service.