11/3/6/91

DOE/ER-0486T

M.L.L.

①

ESnet Program Plan

U.S. Department of Energy
Office of Energy Research
Scientific Computing Staff and
ESnet Steering Committee

March 1991

DOE/ER--0486T

DE91 008242

ESnet Program Plan

U.S. Department of Energy
Office of Energy Research
Scientific Computing Staff and
ESnet Steering Committee
March 1991

MEMORANDUM

Date: March 1991

To: Dr. James F. Decker, Acting Director, Office of Energy Research
Dr. David B. Nelson, Executive Director, Office of Energy Research
Mr. John S. Cavallini, Acting Director, Scientific Computing Staff, Office of Energy Research

The Energy Sciences Network (ESnet) Steering Committee is pleased to submit to you this ESnet Program Plan. This Plan was prepared and submitted at the request of Mr. John S. Cavallini, Acting Director, Scientific Computing Staff, Office of Energy Research. This Plan represents the vision and requirements of the ESnet Steering Committee regarding the future of ESnet.

Within this document you will find a portfolio of information regarding Energy Research networking activities. We have documented the management structure and current ESnet implementation. We have reviewed and represented the networking requirements of the Energy Research programs. We have included a projection of future needs. Appendices containing complementary information can also be found herein.

Given the status of the High Performance Computing and Communications Initiative and related interagency Federal activities, as well as recently introduced legislation, we hope this document will provide you with timely and useful information in that regard.

The ESnet Steering Committee is also taking this opportunity to convey to you our appreciation for the priority and visibility that scientific networking has received under your joint leadership.

On behalf of the ESnet Steering Committee

Sandy Merola

Sandy Merola, Chair

Authors (Energy Sciences Network Steering Committee)

Mr. Robert J. Aiken, *ex officio member*
Scientific Computing Staff, U.S. Department of Energy
Washington, DC 20585

Dr. Raymond A. Bair, *Basic Energy Sciences*
Molecular Science Research Center, Pacific Northwest Laboratories
Richland, WA 99352

Dr. George Brandenburg, *High Energy Physics*
High Energy Physics Laboratory, Harvard University
Cambridge, MA 02138

Mr. Greg Chartrand, *Superconducting Super Collider*
Superconducting Super Collider Laboratory
Dallas, TX 75237

Dr. James Davenport, Basic Energy Sciences
Department of Physics, Brookhaven National Laboratory
Upton, NY 11973

Dr. Martin Greenwald, *Magnetic Fusion Energy*
Plasma Fusion Center, Massachusetts Institute of Technology
Cambridge, MA 02139

Dr. Martha Hoehn, *Nuclear Physics*
Medium Energy Physics Division, Los Alamos National Laboratory
Los Alamos, NM 87545

Dr. Jean-Noel Leboeuf, *Magnetic Fusion Energy*
Fusion Energy Division, Oak Ridge National Laboratory
Oak Ridge, TN 37831

Mr. James Leighton, *ESnet Manager*
National Energy Research Supercomputer Center
Lawrence Livermore National Laboratory
Livermore, CA 94550

Mr. Sandy Merola (Chair), *Advanced Mathematical Sciences*
Information and Computing Sciences Division, Lawrence Berkeley Laboratory
Berkeley, CA 94720

Dr. Larry Price, *High Energy Physics*
High Energy Physics Division, Argonne National Laboratory
Argonne, IL 60439

Dr. Roy Whitney, *ex officio member*
Computer Center Director, Continuous Electron Beam Accelerator Facility
Newport News, VA 23606

Credits

This document was produced at the request of the Scientific Computing Staff, Office of Energy Research, of the U.S. Department of Energy. Contributions were received from Argonne National Laboratory, Brookhaven National Laboratory, the Continuous Electron Beam Accelerator Facility, Lawrence Berkeley Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, Pacific Northwest Laboratory, the Superconducting Super Collider, Harvard University (the High Energy Physics Laboratory), and the Massachusetts Institute of Technology (the Plasma Fusion Center).

Editing and production support was provided by the Information Resources Department of the Information and Computing Sciences Division of Lawrence Berkeley Laboratory.

Contents

Executive Summary

Executive Summary

This Energy Sciences Network (ESnet) Program Plan represents the vision and recommendations of the ESnet Steering Committee regarding ESnet directions and implementations.

The Energy Sciences Network is a nationwide computer data communications network supporting the multiprogram scientific research conducted under the auspices of the Department of Energy (DOE), Office of Energy Research (ER). ESnet provides remote access to unique scientific facilities, including supercomputers, and serves as the needed infrastructure for widely distributed scientific collaborations. Both the High Energy Physics Network (HEPnet) and the Magnetic Fusion Energy Network (MFEnet) rely extensively on the connectivity and transport services afforded by ESnet. ESnet is provided by the Network Group of the National Energy Research Supercomputer Center at the Lawrence Livermore National Laboratory.

The management of ESnet falls under the purview of the Office of Energy Research Scientific Computing Staff. They have established an organization of committees to ensure representation of the scientific user community and local site technical personnel. The ESnet Steering Committee represents the programmatic interests of the user community. The ESnet Coordinating Committee serves to coordinate the need for distributed network management. Task forces and subcommittees are created and dissolved as needed.

The Energy Research community relies heavily on scientific computer networks for all of its current and future activities. ER maintains many unique facilities, large experiments, computer centers, and databases for a geographically dispersed user community. The largest of these facilities include the following:

- the accelerators at the Stanford Linear Accelerator Center (SLAC) and Fermi National Laboratory (FNAL) for the High-Energy Physics program

- the Tokamak Fusion Test Reactor at the Princeton Plasma Physics Laboratory for the Magnetic Fusion Energy program

- various research nuclear reactors at Oak Ridge National Laboratory (ORNL), Brookhaven National Laboratory (BNL), and Argonne National Laboratory (ANL), and photon sources at BNL, Lawrence Berkeley Laboratory (LBL), and ANL for the Basic Energy Sciences program
- the Continuous Electron Beam Accelerator Facility and neutron sources at ANL and Los Alamos National Laboratory (LANL) for the Nuclear Physics program.

Remote access is also required to supercomputers, the National Energy Research Supercomputer Center at Lawrence Livermore National Laboratory (LLNL), for example, and to massively parallel systems at LANL and ANL. Many researchers need access to foreign research facilities, such as the European Center for Nuclear Research (CERN) in Geneva, Switzerland, or the Joint European Torus at Culham Laboratory in England. Databases and computer codes for data analysis, theoretical modeling, and engineering analysis are shared among hundreds or thousands of researchers. In addition to access to the large laboratories and facilities, scientific collaborations are in place between workers at virtually every university and DOE laboratory in the U.S. and between U.S. researchers and their colleagues around the world. Some projects, such as the mapping of the human genome, are carried out jointly by workers at many sites. Excellent communications are an absolute prerequisite for such collaborations. Finally, the next generation of large experimental facilities (projects such as the Superconducting Super Collider, the Compact Ignition Tokamak, and the International Thermonuclear Experimental Reactor) will be designed, built, and operated on a collaborative basis. A new element introduced by these projects is the need for computer links to industrial subcontractors who will be working on design, engineering, and fabrication. These experiments will also involve unprecedented numbers of scientists and will produce huge quantities of data.

The network services required to support these activities range from traditional (electronic mail, file transfer, remote login, remote job entry, remote printing and plotting, bulletin boards) to the more advanced and demanding (distributed file systems, distributed computing with remote procedure calls, video conferencing, and multimedia electronic mail). All the tools now available on local networks will be needed on a national and international scale. For the near future, the network will have to support multiple protocols, such as Internet Protocol (IP), DECnet Phases IV and V, and the International Standards Organization (ISO) protocols, as they are established. Directory services will be required along with gateway access to other scientific networks.

Network bandwidth will need to be increased substantially to accommodate these services. While it is difficult to forecast future demands with accuracy, each of the major ER programs will likely require several tens of megabits per second. Taking all programs together and considering the rate of increase in computer power, the aggregate demand may reach the Gbps level within five years.

Network connectivity will have to be extensive. Every national laboratory and virtually every research university in the country, along with industrial subcontractors and hundreds of foreign sites, will need ESnet service at some level. The highest levels of service would of course be available at sites with the greatest ER activity.

The current ESnet backbone is composed of T1 (1.54 megabits/sec) links, arranged to provide for reliable multipath routing between any two sites, connecting 23 ER sites either directly or indirectly via established mid-level networks. In addition, ESnet is connected with the other major Federal Agency backbones, such as the NSFnet, MILnet, and NASA's NSInet, at two shared internetwork local area networks (LANs) called the Federal Interagency eXchanges (FIXs). One FIX is located at the NASA Ames Laboratory on the west coast, and the other is located at the University of Maryland. Additional interconnects will be established as needed in order to provide for efficient and reliable interagency connectivity.

The number of ER sites connected directly to ESnet is expected to continue to increase. To properly handle the growing number of ER sites and the expanded network requirements of existing ESnet-supported sites and programs, ESnet is scheduled to transition its T1 network to T3 (45 megabits/sec) in the FY 92–93 timeframe.

Introduction 1

Section 1: Introduction

The purpose of the Energy Sciences Network (ESnet) Program Plan is to establish, demonstrate, and represent the vision and recommendations of the ESnet Steering Committee regarding ESnet directions and implementations. This document is intended to assist the Office of Energy Research and the Scientific Computing Staff in ESnet program planning and management, including ESnet prioritization and funding.

This document contains general information about ESnet in terms of both the current and future needs of Energy Research programs for networking infrastructure. Some historical information about Energy Research networking is provided to afford the reader a perspective from which to evaluate the increasing utility of networking to the Energy Research community.

A current ESnet description is provided in detail. ESnet is one of the world's major scientific networks, and the information contained within this document will provide an understanding of ESnet architecture and major components.

ESnet provides the Energy Research community with access to many other peer level networks and a multitude of other interconnected network facilities. The reader will be provided with information concerning ESnet's connectivity and interrelationship to other networks and facilities.

In an attempt to provide adequate background information, summaries of the major Office of Energy Research programs are provided, along with their functional and technical requirements for wide area networking. Major Office of Energy Research programs are managed and coordinated by the Office of Basic Energy Sciences, Office of Fusion Energy, Office of Health and Environmental Research, High Energy and Nuclear Physics, the Superconducting Super Collider, and the Scientific Computing Staff. Trends in the computing environment utilized by these scientific programs have generally resulted in rapidly increasing networking needs. A major section of this document provides information on current and future network utilization as projected and required directly by the Energy Research scientific community.

Networking utilization and forecasting techniques are still in their infancy. The existing complicated architecture, and the integrated yet distributed computing environment, combined with the sophistication of the applications, overwhelm the ability of current network utilization and forecasting tools and techniques. Nevertheless, in addition to the programmatic forecast of networking requirements contained herein, we provide current and future utilization information from a DOE site point of view. Only with the programmatic and site utilization information, combined with a forecast of technology trends and programmatic directions, can network engineers begin to design the network of the future. However, as networking is an enabling technology, future novel uses of the network are impossible to predict; and, as a result, future network utilization will likely exceed that forecasted from known uses. Thus, providing some overcapacity is prudent planning.

There exist certain widespread advances in scientific network technology, mostly in the area of interconnectivity and interoperability. Various standards, standard groups, and collaborations among the government, industry, and university sectors have fostered such advances. Those directly related to Energy Research networking are mentioned in this document.

Major Federal Initiatives, such as the Federal High Performance Computing and Communications (HPCC) Program and the National Research and Educational Network, are providing a national focus and support for recognition of networking as an important enabling technology for the advancement of science and industrial competitiveness throughout the scientific, industrial, and educational communities of the United States. The relationship of the HPCC components is shown in Figure 1-1, taken from Reference 1. Mention of these important developments is included herein.

No program plan would be complete without an accompanying financial budget. This document includes needed financial forecasts underlying the successful growth of the Energy Research network infrastructure.

Reference:

1. *The Federal High Performance Computing Program*, Office of Science and Technology Policy, 1989.

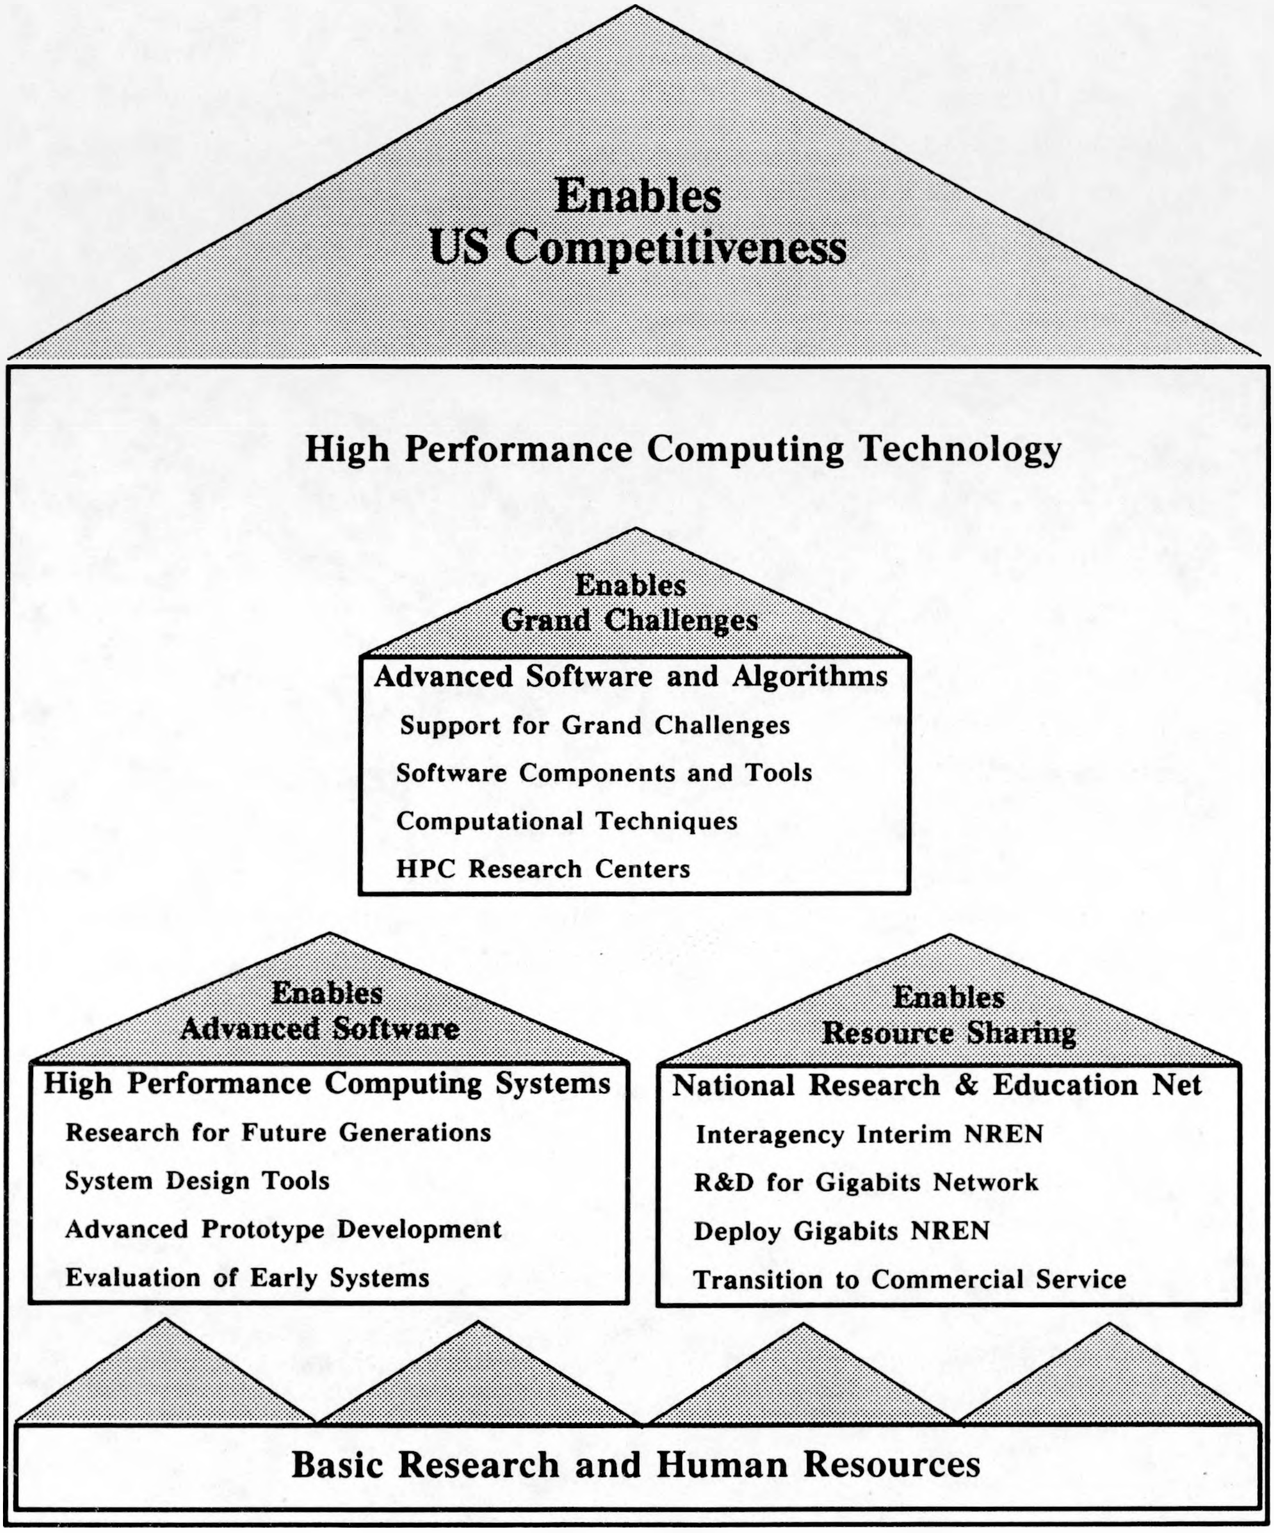

Figure 1-1. Relationship of High Performance Computing and Communications Program components.

Background 2

Section 2: Background

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 supported by ER. These programs include High Energy Physics, Nuclear Physics, Materials Sciences, Chemical Sciences, Carbon Dioxide Research, Engineering and Geosciences, Fusion Energy and Applied Plasma Physics, and Health and Environmental Research. To provide these computational resources, the Scientific Computing Staff manages the following:

- The National Energy Research Supercomputer Center (NERSC), which operates the following supercomputer systems: Cray 2's (three systems), Cray X-MP/22 (one system), and Cray 1S (one system) (see Figure 2-1 and Table 2-1).

- The Supercomputer Computations Research Institute (SCRI) at Florida State University (FSU), where research is carried out in the computational sciences related to ER mission areas. SCRI operates a Cray Y-MP and 65000-node CM-2 (see Figure 2-2).

- ESnet, which is a nationwide data communications network, facilitating remote access to unique Energy Research facilities [for example, the facilities described in *Scientific User Facilities, A National Resource* (see Figure 2-3)], [1] providing needed information dissemination among ER scientific collaborators, and providing widespread access to ER supercomputer facilities.

- Several advanced computational research centers, such as the Advanced Computational Research Facility (ACRF) at Argonne National Laboratory (ANL), which operate research/experimental computational resources. For example, hypercube systems are used 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, the National Aeronautics and Space Administration, the National Cancer Institute, the Department of Defense, and others. As a result of

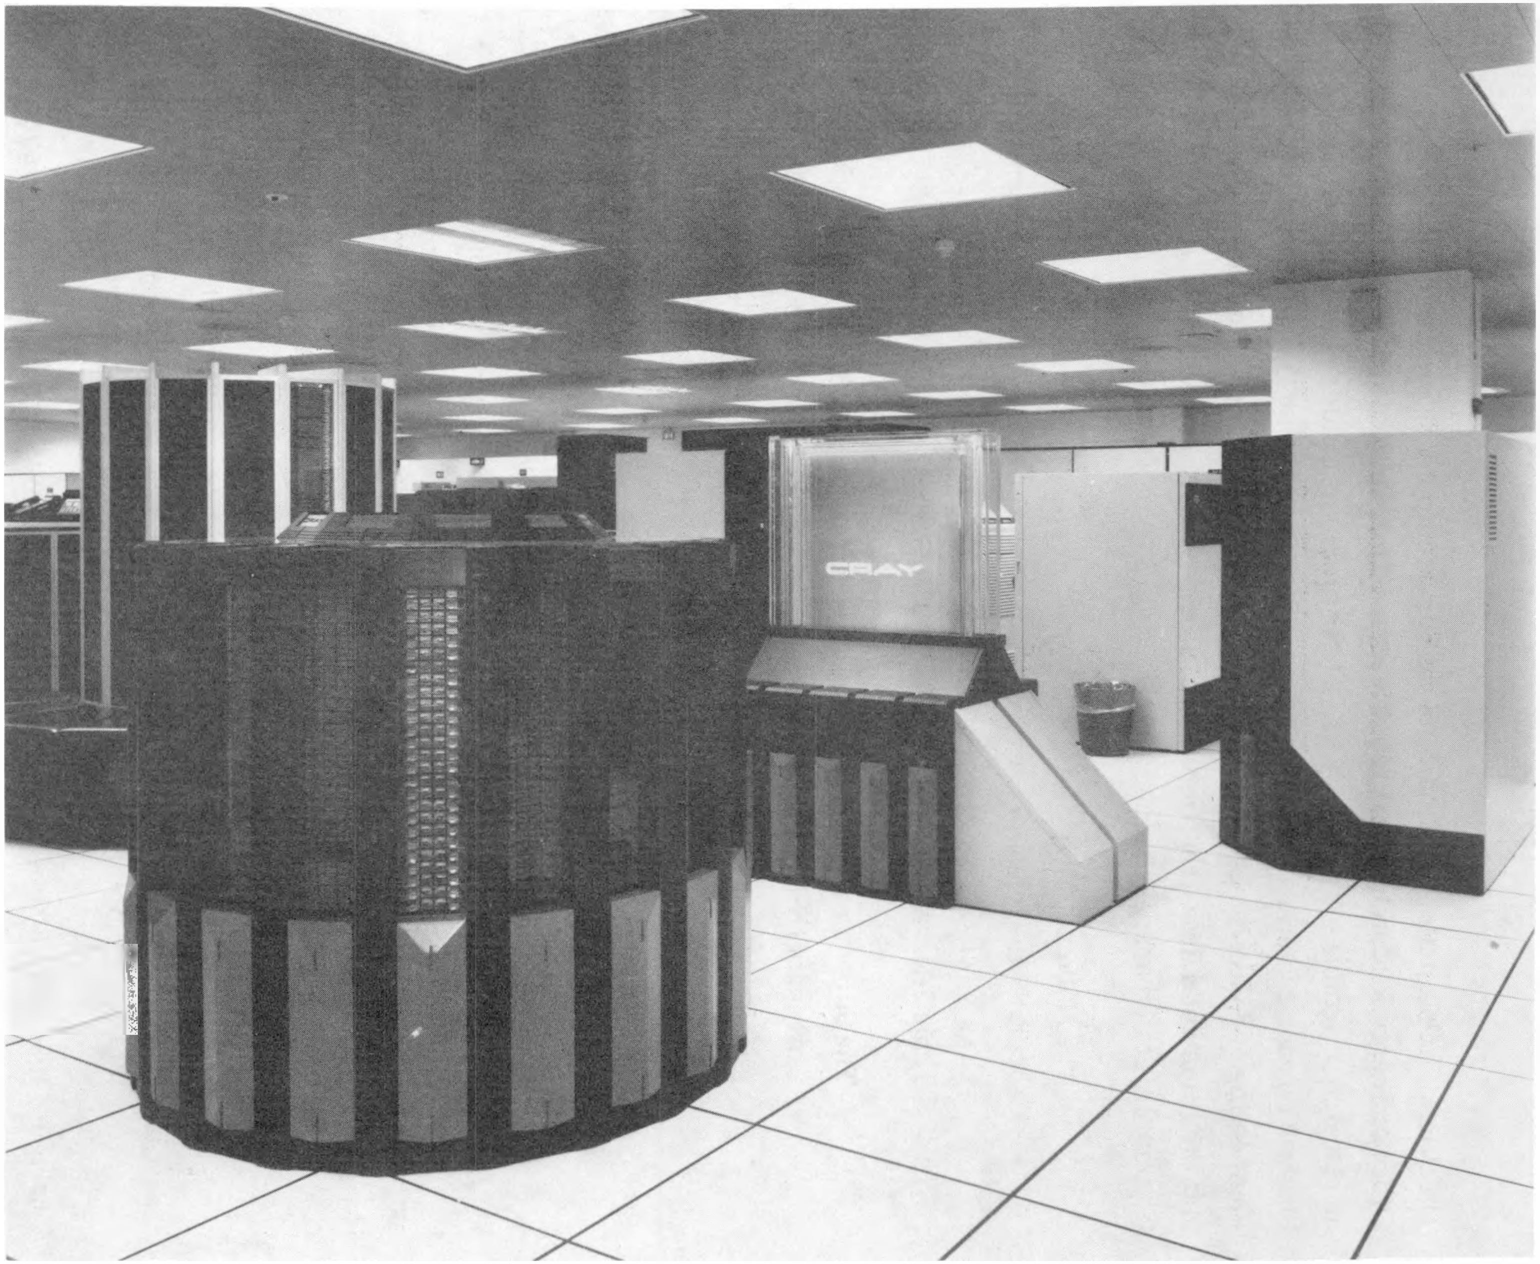

Figure 2-1. NERSC eight-processor Cray-2 supercomputer.

Table 2-1. Applied Mathematical Sciences Supercomputer Resources Allocated (1000 Cray hours).

	FY 1991		FY 1990	
Energy Research Program	Requests	Allocations	Requests	Allocations
High Energy and Nuclear Physics (including SSC)	106.5	59.3	195.1	51.6
Basic Energy Sciences	83.5	40.9	76.7	33.7
Health and Environmental Research	29.5	16.7	20.5	12.9
CHAMMP (Global Climate Change)	36.0	15.0	15.0	12.0
Applied Mathematical Sciences	13.5	5.0	9.6	3.8
Fusion Energy	159.7	75.0	110.8	64.0
Totals	428.7	211.9	427.7	178.0

Distribution of Resources

DOE Laboratories	62%
University	35%
Industry	3%

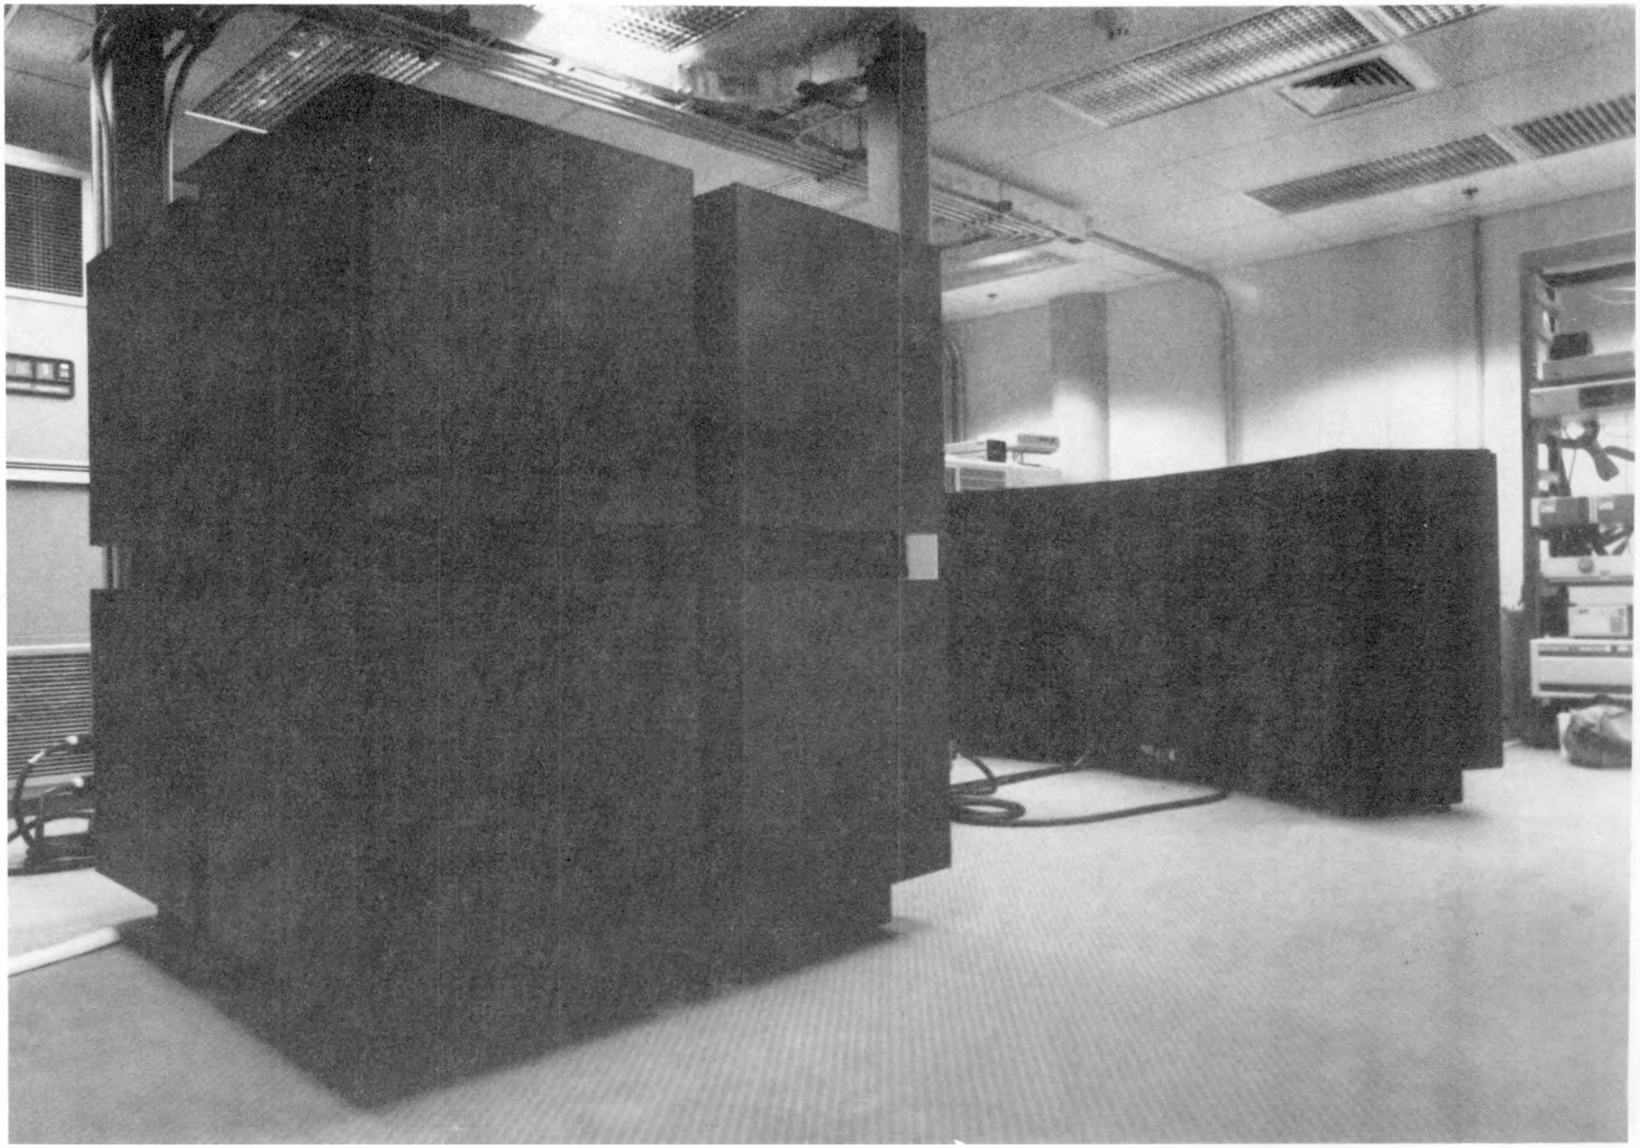

Figure 2-2. "Connection Machine" at Florida State University's Supercomputer Computations Research Institute, manufactured by Thinking Machines, Inc.

Figure 2-3. Scientific user facilities operated by DOE's Office of Basic Energy Sciences.

these initiatives and in recognition of its importance to scientific collaboration support, computer networking emerged as an important underlying and enabling technology.

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. MFEnet, which was developed in 1975, continued to provide the fusion scientific community with dedicated network access, including electronic mail, file transfer, and terminal access to the National Magnetic Fusion Energy Computer Center. Some 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 (HEP), a private network of leased 9600-baud lines running DECnet grew as university groups required connectivity between their local minicomputers and the facilities at the Stanford Linear Accelerator Center, Fermi National Accelerator Laboratory, Brookhaven National Laboratory, and other HEP laboratories.

In 1985, a subpanel of the High Energy Physics Advisory Panel (HEPAP), chaired by J. Ballam, produced 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 [2] 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.

As a result, a general-purpose computer-network architectural concept was developed for ER cross-program computer networking support. The redesign of the existing MFEnet, the Energy Research supercomputer access network, was evaluated for its ability to serve as the infrastructure for the modernization of needed ER networking. This concept was reviewed by an interagency working group in January 1986 and was determined to be a major step toward the establishment of an interagency internet capability. This internet architectural model (see Figure 2-4) 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 other sponsored national and regional scientific computer networks.

In response to Dr. Trivelpiece's charge, the Scientific Computing Staff proposed an ESnet concept that

- recommended the formation of the Energy Sciences Network Steering Committee to represent the ER scientific community

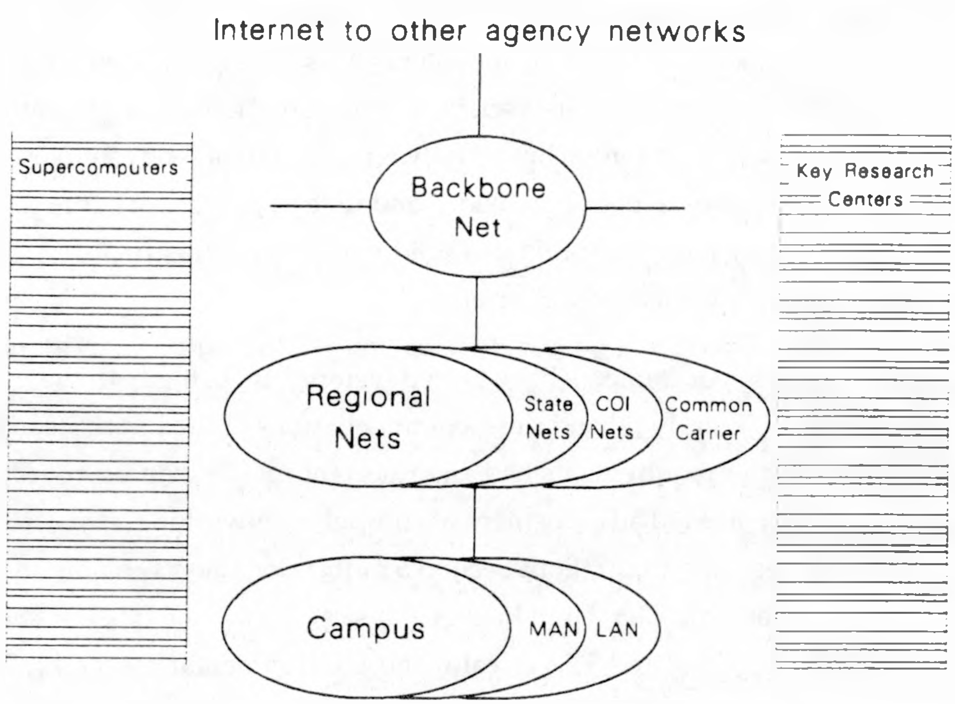

Figure 2-4. Research network hierarchy as of 1989.

- offered the new MFEnet design and facilities for consideration as the basis for ESnet facilities

- proposed an evolutionary operational model to guide the needed network transition and upgrade, 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 United States House of Representatives Science and Technology 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, and Congress mandated a computer network study in the FY 1987 NSF authorization bill. Members of the ESnet Steering Committee have made significant contributions in the aforementioned areas. The ER computer network committees, such as the ESnet Steering Committee (ESSC), are also aware of the national focus and scope of the National Research and Education Network (NREN) and are formulating appropriate plans for ER's participation in NREN as it evolves. This has significantly impacted government- and university-sponsored scientific computer networking throughout the United States. Bandwidth and interconnectivity among networked resources and collaborators has increased in both bandwidth and ubiquity. Network connectivity has transformed from being a useful option to a solid part of the infrastructure underpinning the advancement of U.S. science.

ER computer network requirements have been developed by using both the Congressional FCCSET study [2] and the recommendations of the ESnet Steering Committee. The former requirements had an important role in shaping government-wide planning for the needed advancement of computer networking. The latter requirements are presented throughout Section 5 of this document, notably in Sections 5.2 (Basic Energy Sciences), 5.3.3 (High Energy Physics), 5.4.3–5.4.5 (Fusion Energy), 5.5.2 (Nuclear Physics), 5.6.7 and 5.6.9 (Health and Environmental Research), 5.7.2 (Scientific Computing), and 5.8.2 (the Superconducting Super Collider). A summary of future network requirements is described in Section 7. Also presented in this document are a summary of the ESnet Steering Committee activities (Section 3), a statement of the benefits anticipated from the ESnet computer networking project (throughout Section 5), and a plan of action, including budget requirements through FY 1996 (Section 7).

References:

1. *Scientific User Facilities, A National Resource,* DOE/BES publication.

2. D.F. Stevens, "DOE Networking Requirements," in *A Report to the Office of Science and Technology Policy on Computer Networks to Support Research in the United States, Vol. III: A Compendium of Supporting Technical Data,* Federal Coordinating Committee on Science, Engineering, and Technology, June 1987, pp. 169–178.

ESnet Management Structure 3

Section 3: ESnet Management Structure

3.1 Introduction

This section provides information about the management and coordination activities associated with the planning and operation of the Energy Sciences Network (ESnet). ESnet is sponsored by the Scientific Computing Staff (SCS) of the Office of Energy Research of the U.S. Department of Energy. SCS is responsible for the funding of the network and oversees its management.

In order to ensure the success of ESnet, a tiered committee structure has been implemented. The Network Access Committee (NAC) is composed of ER program managers who meet to assist SCS in budgetary planning and provide programmatic input concerning network policy and planning. The ESnet Steering Committee (ESSC) represents the programmatic interest from the user's perspective to the SCS on network policy and funding. The technical management of the ESnet backbone is the responsibility of the ESnet Manager and Staff (ESMAN) at Lawrence Livermore National Laboratory. The ESnet Site Coordinating Committee (ESCC) reports to ESSC and advises the ESnet Manager on network management issues. ESCC establishes task forces, such as the ESnet DECnet Working Group (EDWG), to deal with specific technical issues. The interrelationship of these groups is diagrammed in Figure 3-1. Following Figure 3-1 are the charters, method of selection of leadership, and duration of the various entities.

3.2 Scientific Computing Staff (SCS)

The SCS has total programmatic management responsibility for ESnet. The ESnet program manager within SCS is designated by the Director of the SCS.

3.3 Network Access Committee (NAC)

The NAC, composed of representatives of the ER program offices, was established in 1990. Its charter is still under discussion, but a preliminary charter is the following: "The Network Access Committee (NAC) shall be composed of one member from each of the SSC, HENP, BES, AMS, HER and FE program offices. Their function will be to review implementation plans of the research networks associated with ER with respect to

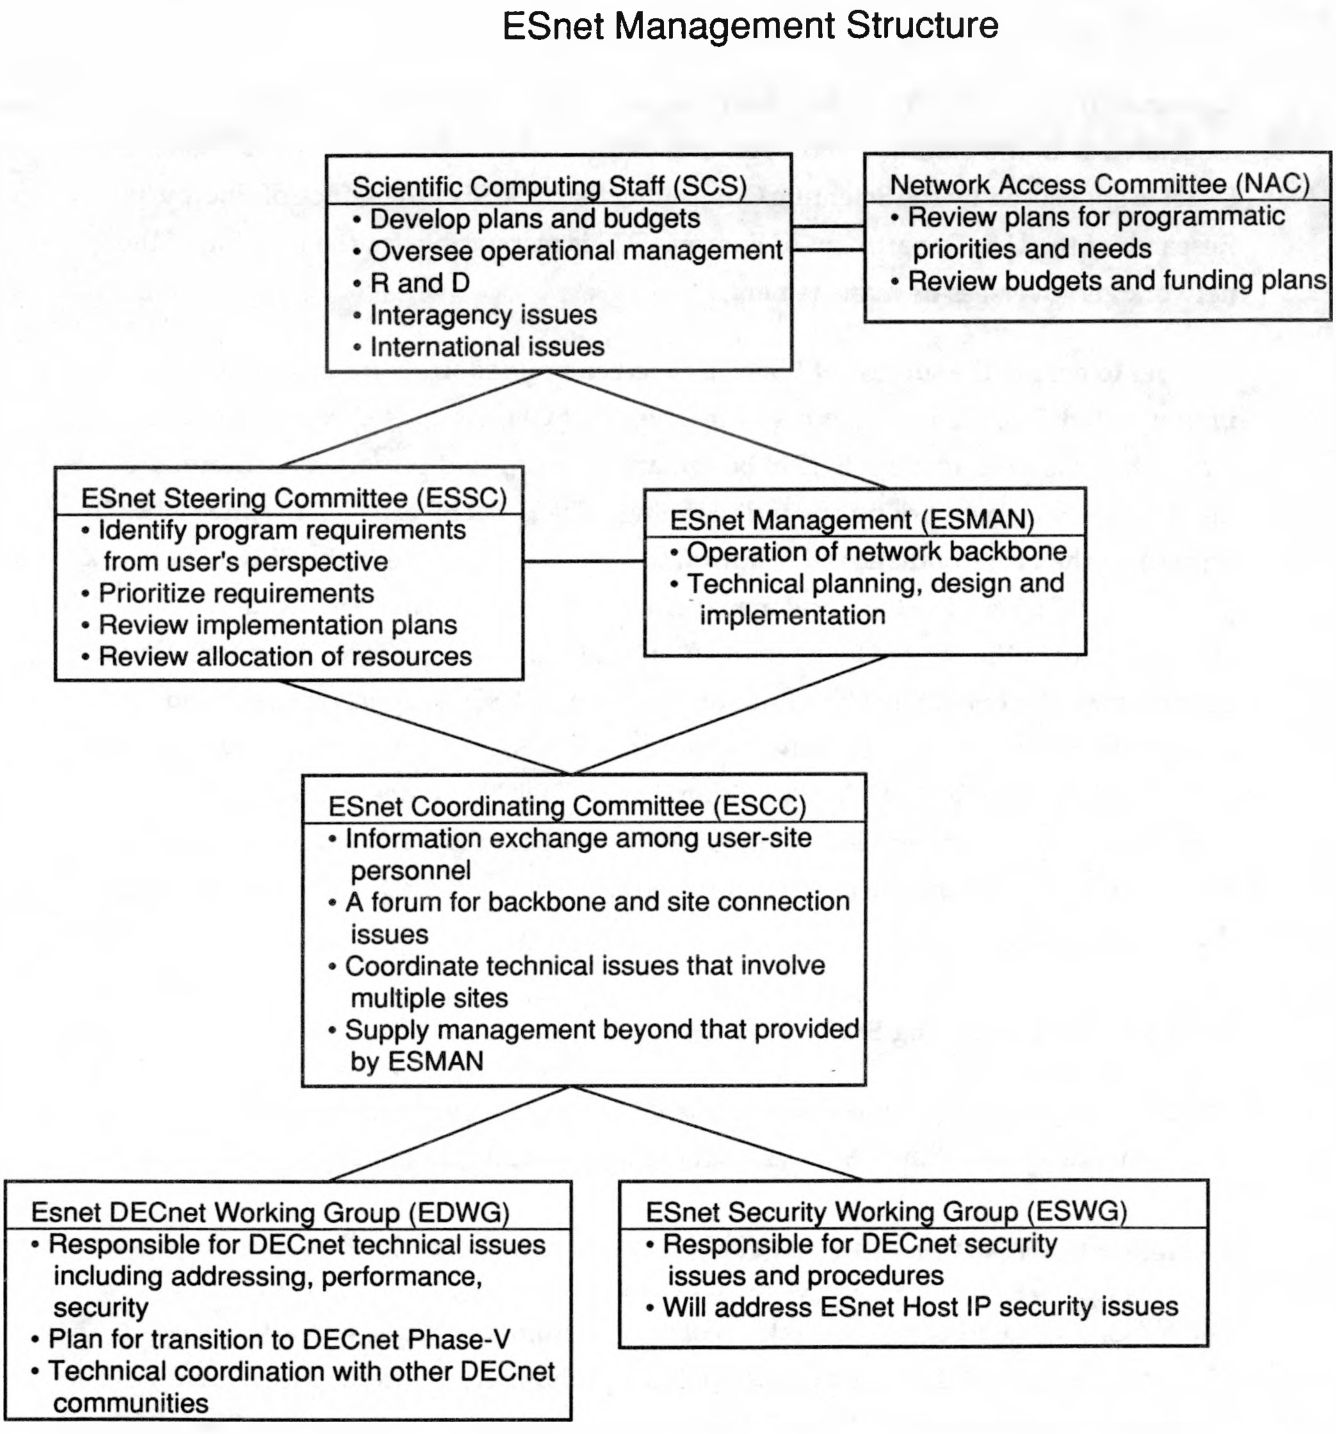

Figure 3-1. ESnet management structure.

programmatic needs and priorities. In addition, the NAC should assist in determining the future budgetary requirements for ER funds that are necessary for the networks to achieve the research mission of DOE." Each member of the NAC is appointed by the Office Director of the ER program he or she represents. The member from SCS serves as chair. The NAC is a standing committee.

3.4 ESnet Steering Committee (ESSC)

The ESSC, composed of representatives of all ER programs, was established in 1986 and is charged to:

- document, review, and prioritize network requirements for all ER programs
- ensure that ESnet goals are achieved without impacting ongoing program requirements
- establish ESnet computer network performance objectives
- propose innovative techniques for enhancing ESnet capabilities
- advise the ESnet manager and staff
- identify research needs for addressing network requirements
- review the ESnet budget as presented by the Scientific Computing Staff with respect to the prioritized network requirements.

The members of ESSC are appointed by the Division Director of the ER program he or she represents. Its chairman is appointed by the SCS from within the committee with the advice and consent of the committee. ESSC is a standing committee.

The High Performance Computing and Communications (HPCC) Program and the National Research and Education Network (NREN) initiatives are of the utmost importance to the DOE ER programs. These initiatives directly influence the ability of government agencies to provide the needed high-performance computing environment, including networking to the U.S. Government–sponsored scientific research community. Their passage and implementation require that the ESnet committees assist with the definition of NREN, thereby ensuring that DOE's mission-driven program requirements are met. In this vein, ESSC is developing the requirements for the relationship between the DOE ER programs/networks and NREN. ESSC will ensure that members of the various ESnet committees (ESSC, ESCC, EDWG) will be part of the requisite NREN steering and coordinating committees, the NREN user group, and the Federal Engineering Planning Group (FEPG).

Through this representation the DOE community will help steer NREN in a direction that guarantees the continuance of the DOE mission and the fulfillment of its program requirements. DOE maintains a strong and broad programmatic presence by providing sizable funding and support for numerous researchers at many universities, colleges, and laboratories. Therefore, DOE expects to have a measurable impact on the management, design, implementation, funding, and operation of NREN.

3.5 ESnet Site Coordinating Committee (ESCC)

ESCC, composed of representatives from each of the major ESnet backbone sites, was established in 1987 to serve as:

- an advising body to the ESnet Steering Committee on the broad range of technical issues related to Energy Research scientific computer networks
- a forum for information interchange on ESnet activities and plans and site needs and plans
- a forum for interaction with the ESnet manager and staff
- a forum for consideration of technical networking issues connected to ESnet for which coordinated multiple-site involvement is necessary or advantageous
- a forum for interacting with ER program-specific networking activities that use or would like to use ESnet facilities
- a body to supply complementary network management to the ESnet manager and staff.

The members of ESCC are appointed by the individual ESnet site organizations. Current membership represents 22 interconnected Energy Research sites. The chairman of ESCC is appointed by SCS from within the committee, with the advice and consent of both

ESSC and ESCC. To carry out the above functions, ESCC will appoint various working groups and task forces as the need arises. A working group will exist for an extended period of time to address issues within a general category (e.g., ESnet/DECnet issues, security issues). A task force will exist for a short term to accomplish a narrow, well defined goal (e.g., "DECnet routing," "TCP/IP routing"). The membership of working groups is subject to approval by ESCC. Working groups choose their own chair, subject to approval by ESCC. Task forces may have a more flexible structure depending on the issue, but in any case the leadership and membership is subject to the approval of ESCC. ESCC is a standing committee. The ESCC working groups are discussed in Section 6 of this document.

Current Status 4

Section 4: Current Status

4.1 Introduction

The Energy Sciences Network (ESnet) is a nationwide computer data communications network managed and funded by the U.S. Department of Energy, Office of Energy Research (DOE/OER), for the purpose of supporting multiple program, open scientific research. ESnet is intended to facilitate remote access to major Energy Research (ER) scientific facilities, provide needed information dissemination among scientific collaborators throughout all ER programs, and provide widespread access to supercomputer facilities.

ESnet is engineered, installed, and operated by the networking staff of the National Energy Research Supercomputer Center (NERSC), located at Lawrence Livermore National Laboratory.

Details of the ESnet usage guidelines and procedures, including security, are described in Appendix B.

4.2 Description

The ESnet 1991 T1 backbone topology is shown in Figure 4-1.

ESnet began operational deployment of a T1 (1.3 to 1.5 Mbps) circuit-based backbone in late 1989. It became fully operational with the initial configuration in early 1990 and currently has 22 major OER-supported sites directly connected to the backbone.

The majority of the T1 circuits are supplied by AT&T through the GSA FTS-2000 program. These T1 circuits are all fiber-optic based, and all support Extended Superframe Format (ESF), allowing significant diagnostic monitoring while the link is in use. Additionally, many of the circuits have "clear channel" capability, allowing use of the full 1.536-Mbps bandwidth for user data.

The balance of the T1 circuits are provided through a collaborative effort with the Defense Advanced Research Projects Agency (DARPA), as part of the National Network Testbed (NNT) project. These circuits are also generally fiber-optic and include ESF but do not support clear channel capability, and therefore provide 1.344-Mbps user data bandwidth.

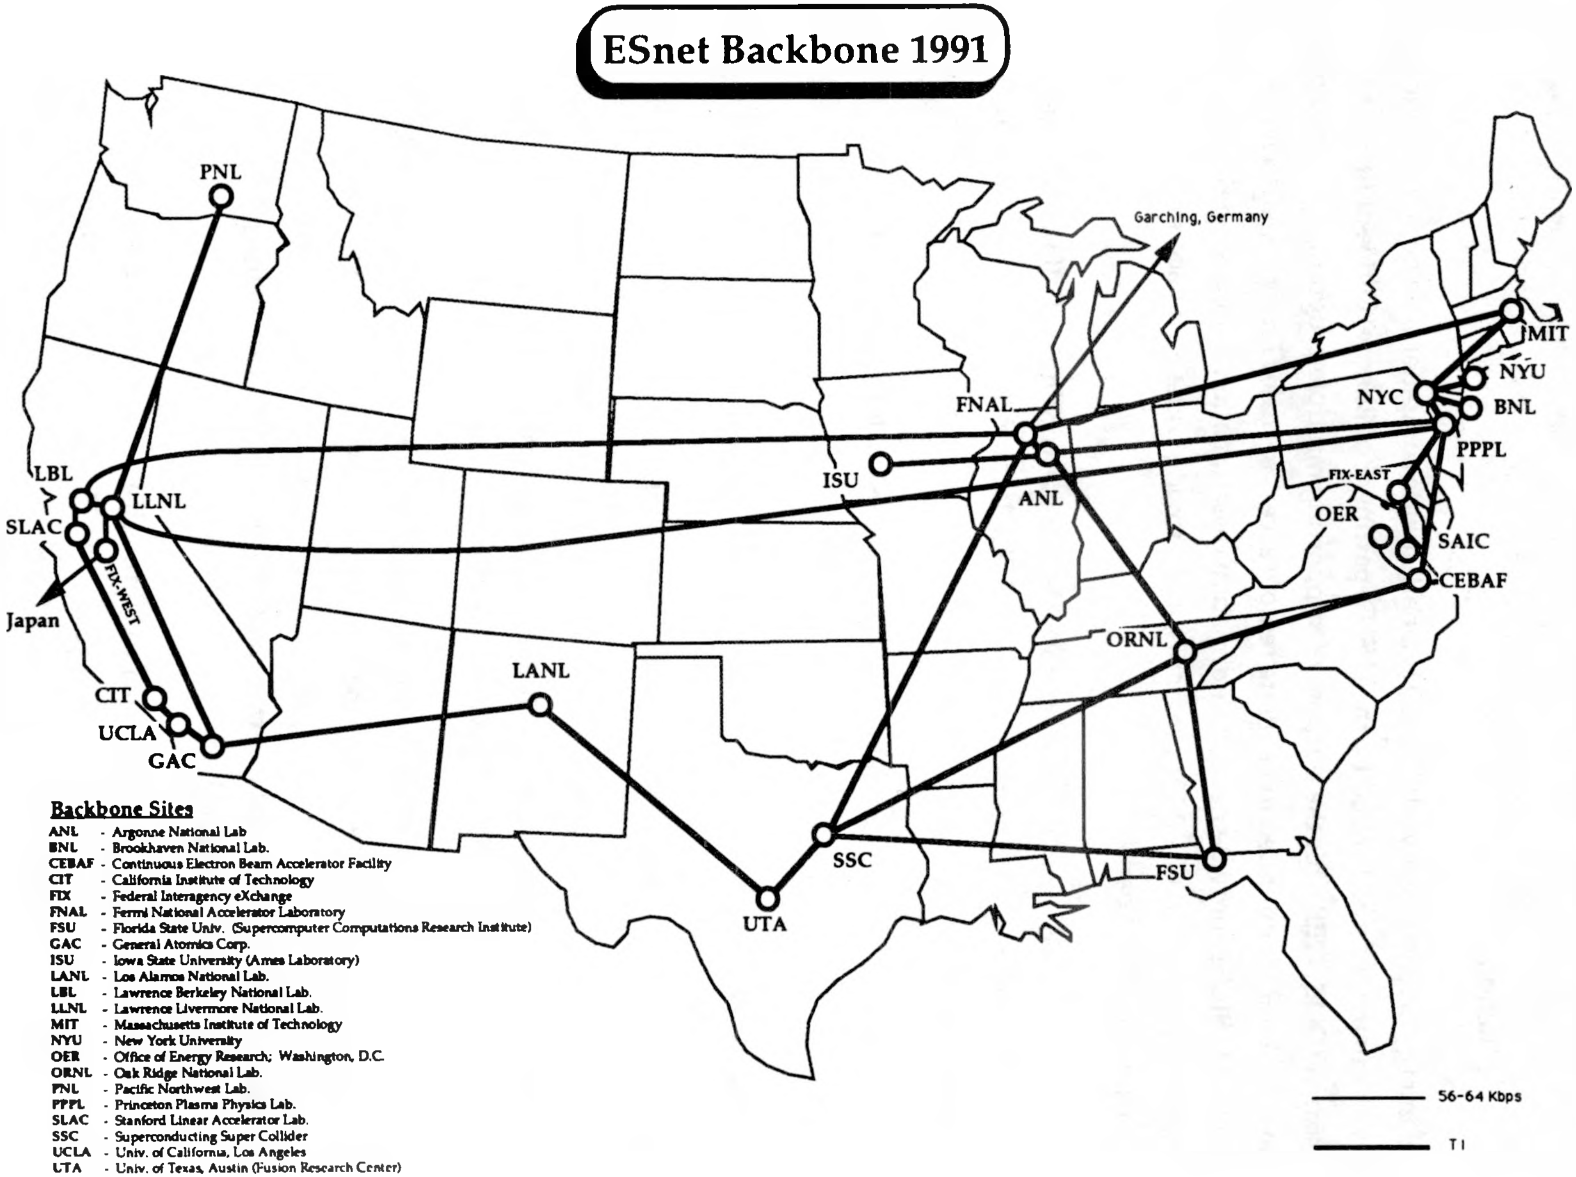

Figure 4-1. ESnet backbone topology, 1991.

Network routers currently in use are products of cisco Systems, Inc.

Protocols supported are Department of Defense Internet Protocol (DOD-IP) and DECnet Phase IV. Switching of X.25 packets was planned to be added in late 1990. Capability of supporting the Open Standards Interconnect (OSI) Connectionless Network Protocol (CLNP) is planned for mid-1991. Upgrades to bandwidths beyond T1 speeds (1.544 Mbps) are under active consideration.

4.3 What Does it Connect?

The implementation of internet protocol (IP) routing on ESnet is intended to allow backbone sites to interoperate with each other and with other non-OER resources, but not necessarily to allow the non-OER resources to interoperate over ESnet. For example, regional or mid-level networks may communicate with backbone sites over ESnet, but the same regionals are not allowed to interoperate with each other via ESnet. Figure 4-2 shows the current Internet structure.

4.3.1 Backbone Sites

The following is a list of sites that are directly connected to the ESnet T1 backbone as of November 1990:

1. Argonne National Laboratory
2. Brookhaven National Laboratory
3. Continuous Electron Beam Accelerator Facility
4. California Institute of Technology
5. DOE, Office of Energy Research
6. Fermi National Accelerator Laboratory
7. Florida State University
8. General Atomics
9. Iowa State University
10. Los Alamos National Laboratory
11. Lawrence Berkeley Laboratory
12. Lawrence Livermore National Laboratory
13. Massachusetts Institute of Technology
14. New York University
15. Oak Ridge National Laboratory

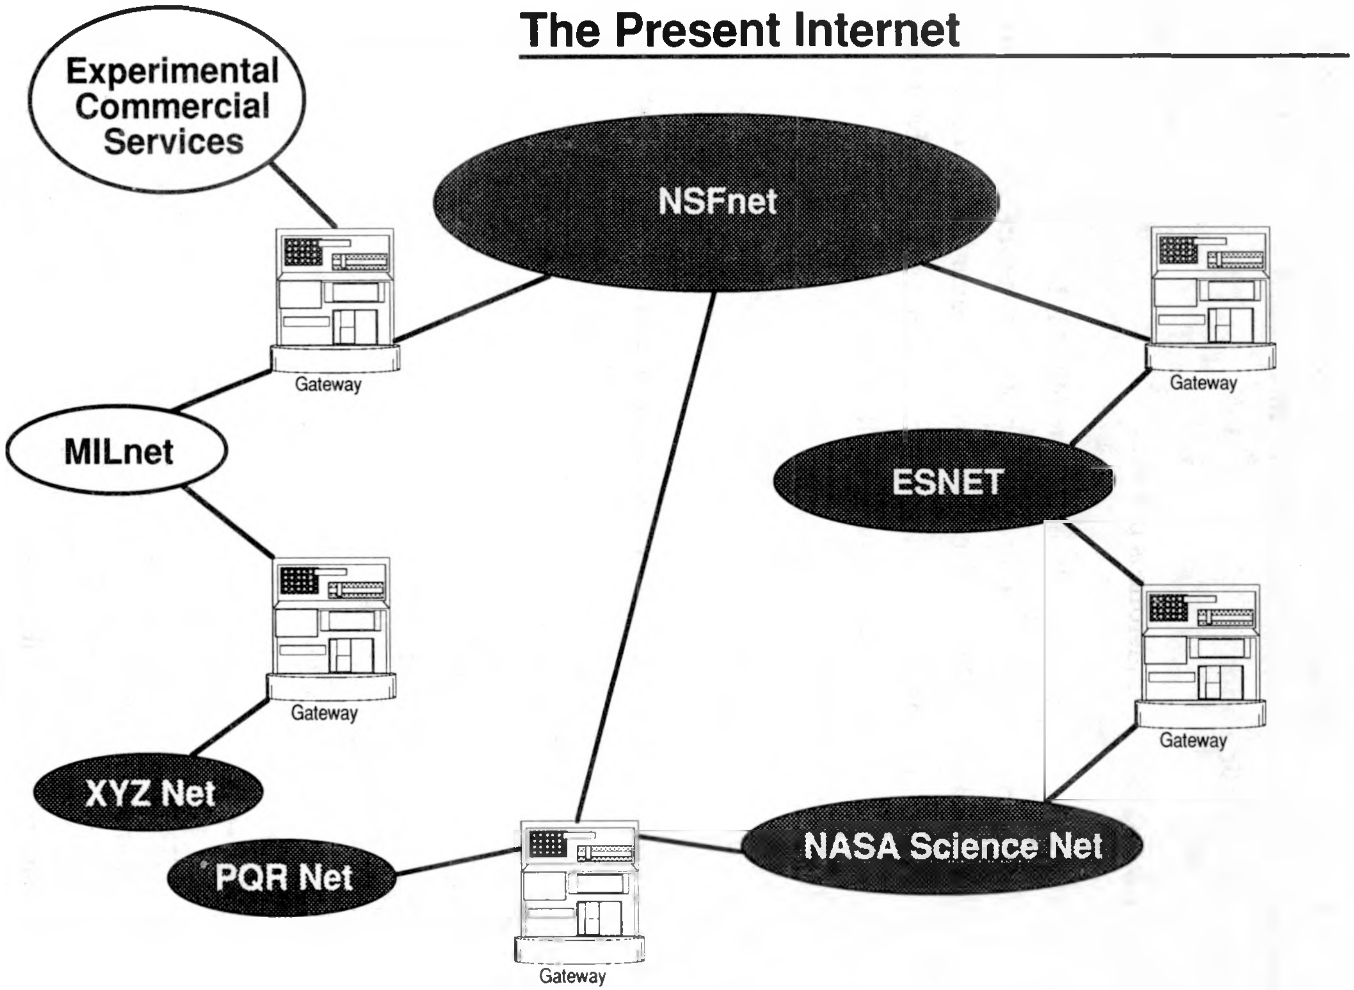

Figure 4-2. Current Internet Structure

16. Pacific Northwest Laboratory
17. Princeton Plasma Physics Laboratory
18. Scientific Applications, Inc.
19. Stanford Linear Accelerator Center
20. Superconducting Super Collider
21. University of California at Los Angeles
22. University of Texas at Austin

Sandia National Laboratory at Albuquerque is planned to be added during during early 1991.

4.3.2 Other National Networks

ESnet currently has two Federal Interagency eXchange (FIX) connections established to allow exchange of data with the Defense Communications Agency (MILnet), NASA (NASA Science Network, or NSN), and NSF (NSFnet) (see Figure 4-3). The FIX-West interconnect is located at the NASA Ames Research Center, Mountain View, California; and FIX-East is located near the University of Maryland.

The interconnect with NSN supports DECnet Phase IV traffic, as well as IP.

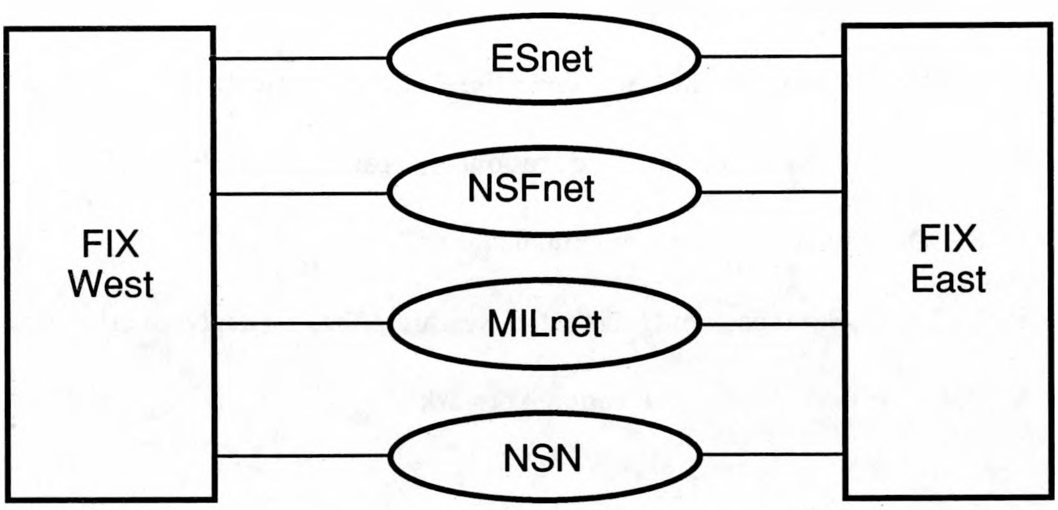

Figure 4-3. FIX connections.

4.3.3 *International Networks*

New international connectivity is planned to both Japan and Germany during 1990 and 1991.

The connection to Japan will use a shared fiber-optic 512-Kbps trunk to Hawaii, and several 64-Kbps links to Japan from Hawaii. Access to many sites within Japan is then available via two Japanese internal networks, TISN and WIDE.

Similarly, access to Germany will be via a shared 128-Kbps fiber-optic trunk between PPPL (Princeton Plasma Physics Laboratory, Princeton, NJ) and Bonn, Germany. Access within Germany will then be provided to nearly all major research facilities and universities through an internal X.25-based network, WIN.

4.3.4 *Regional Networks*

The following list shows the regional networks with which ESnet is currently interconnected. Figure 4-4 illustrates this interconnection.

- BARRnet: San Francisco Bay Area Regional Research Network
- CERFnet: California Education and Research Federation Network
- Los Nettos: Los Angeles Regional Network
- NEARnet: New England Academic and Research Network
- NYSERnet: New York State Educational Research Network
- SESQUInet: Texas Sesquicentennial Network
- SURAnet: Southeastern Universities Research Association Network
- THEnet: Texas Higher Education Network

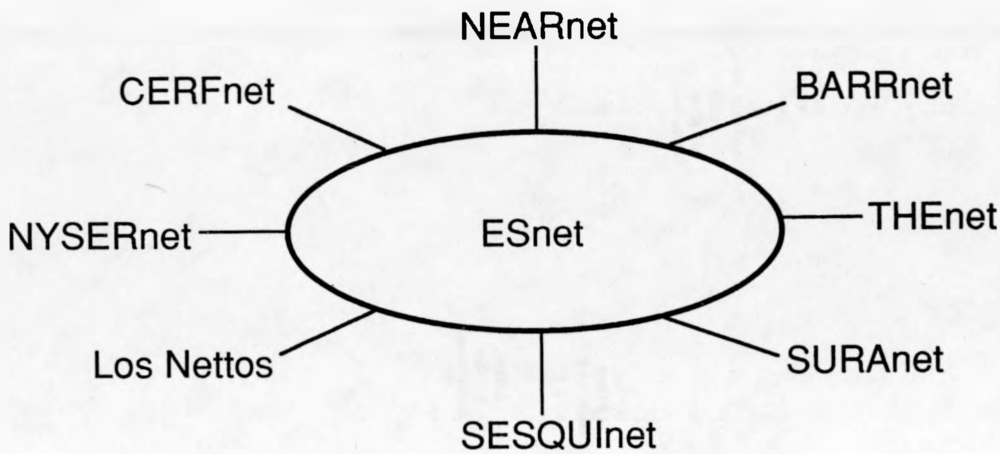

Figure 4-4. Regional networks interconnected to ESnet.

The following is a list of additional regional networks that will be interconnected with ESnet during 1991:

- CICnet: Committee on Institutional Cooperation Network
- MIDnet: Midwestern States Network
- NORTHWESTnet: Northwestern States Network
- WESTnet: Southwestern States Network

4.3.5 HEPnet

The ER High Energy Physics (HEP) research community has created HEPnet to support that community's networking needs. The primary purpose of this network is to facilitate the geographically dispersed collaborations typical of HEP research projects. The ESnet backbone is used by HEPnet for high-speed interconnections. Figures 4-5 through 4-7 show the major parts of HEPnet.

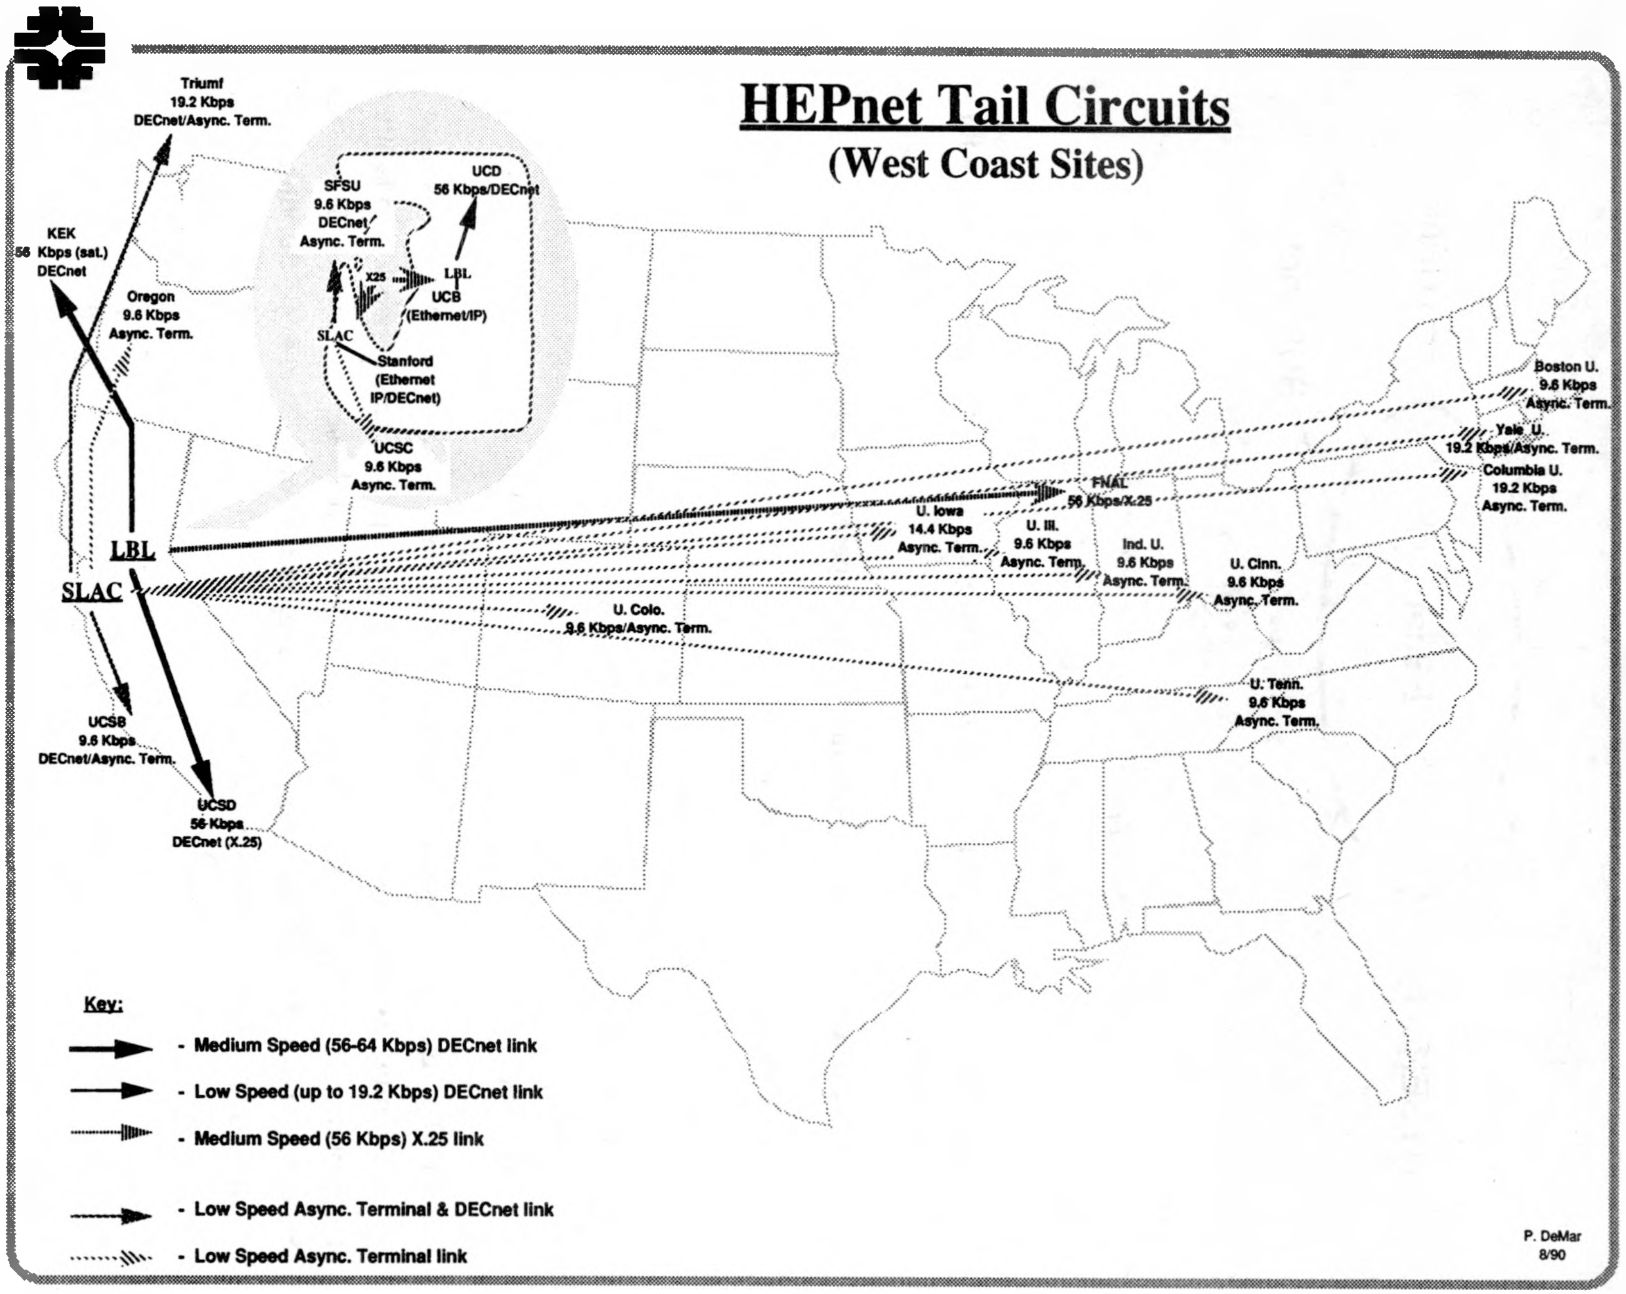

Figure 4-5. HEPnet tail circuits (West Coast sites).

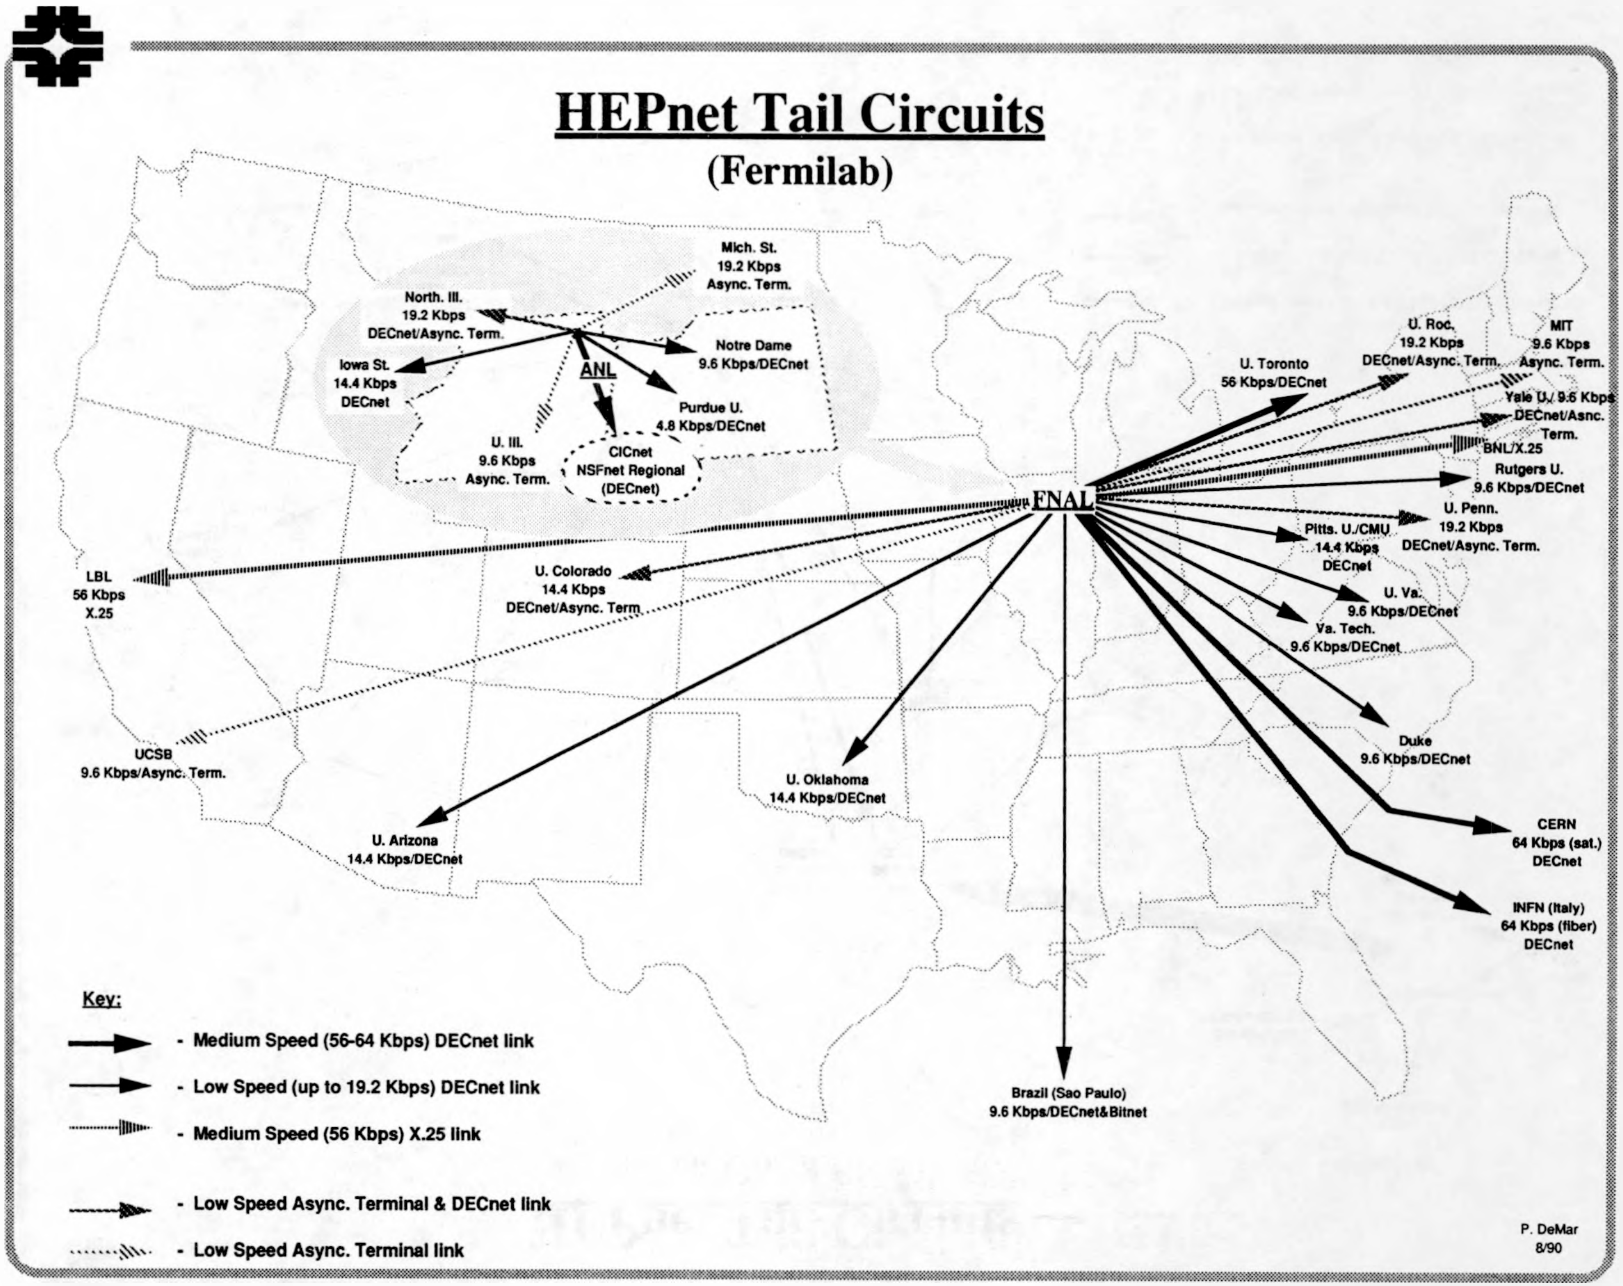

Figure 4-6. HEPnet tail circuits (Fermilab).

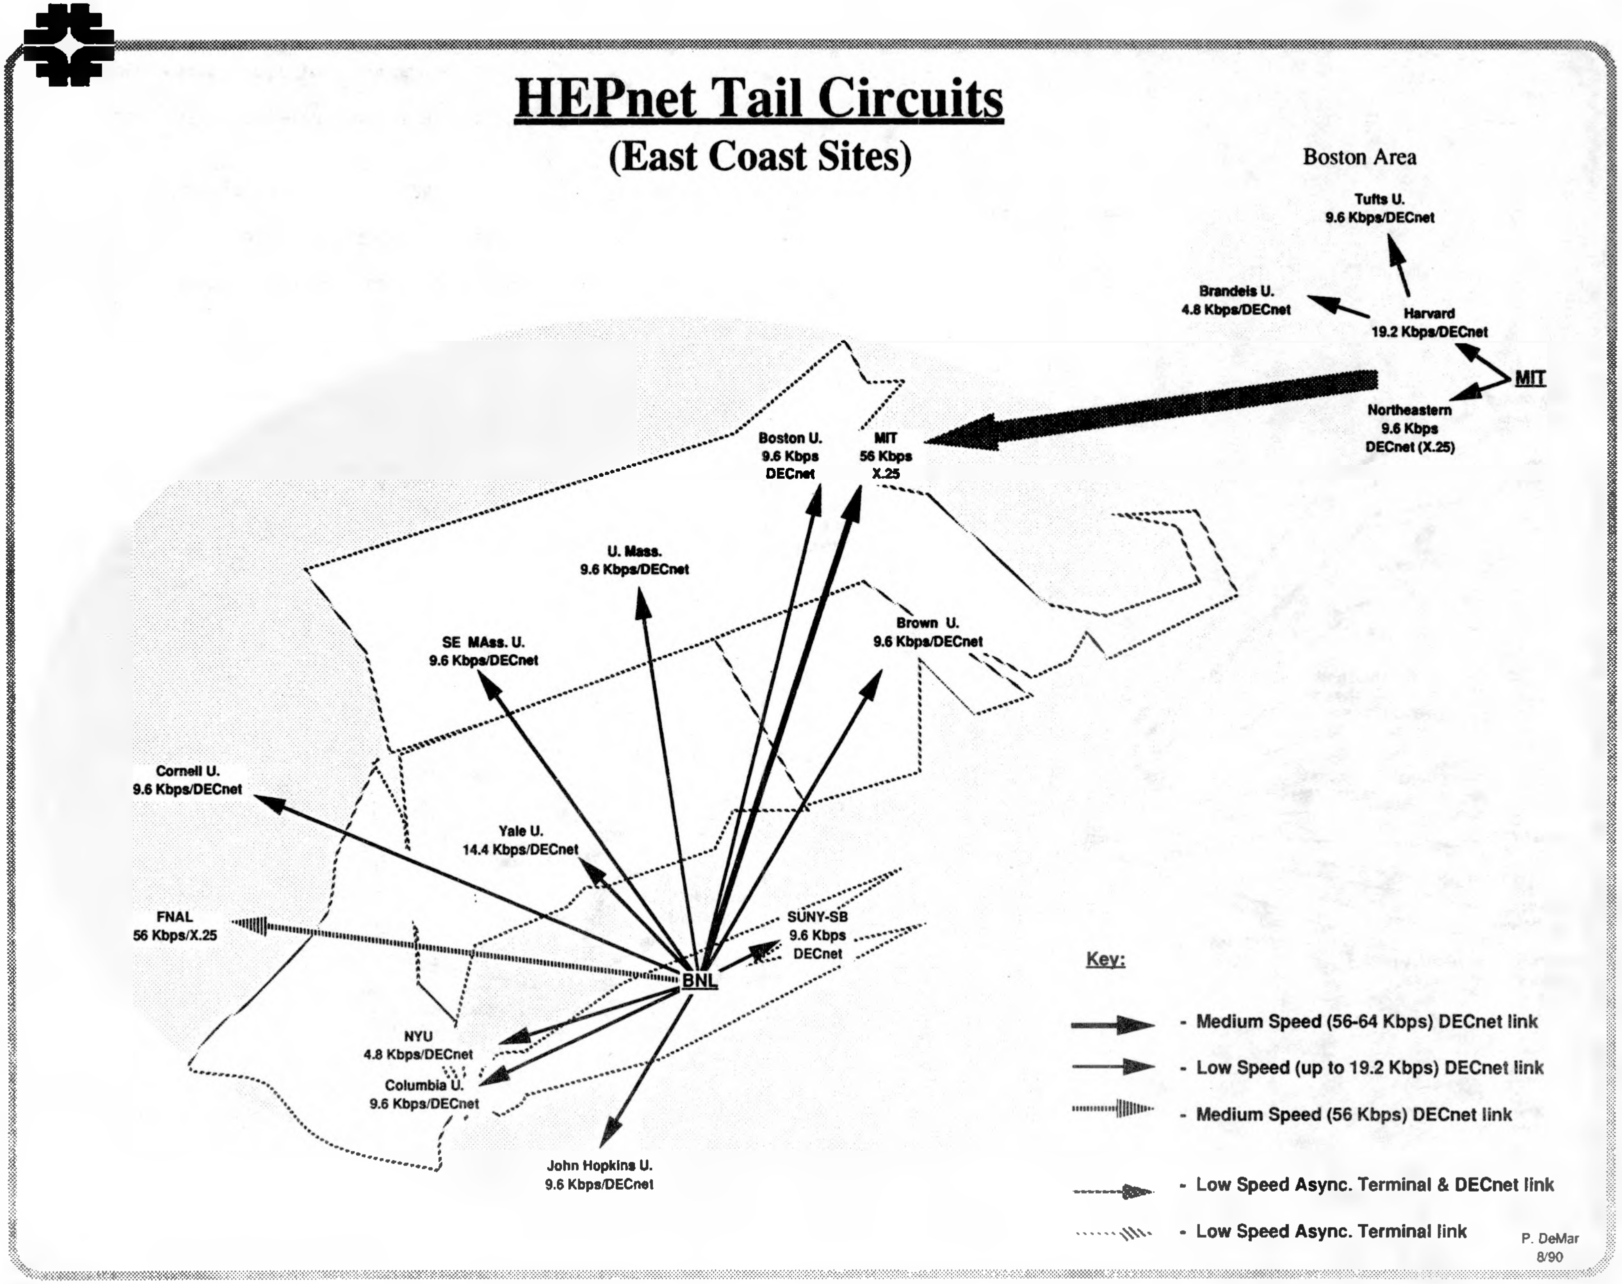

Figure 4-7. HEPnet tail circuits (East Coast sites).

4.3.6 *MFEnet*

MFEnet (Magnetic Fusion Energy Network) was created in 1976 to provide access to the National Magnetic Fusion Energy Computer Center (NMFECC, since renamed NERSC), located at Lawrence Livermore National Laboratory. MFEnet uses its own protocol family, layered on top of Internet IP. MFEnet is implemented for VAX/VMS and the Cray timesharing system (CTSS) operating system used on the NERSC Crays. The ESnet backbone serves as a carrier network for MFEnet. Figure 4-8 shows MFEnet as of 1991.

4.3.7 *Application Layer Gateways*

Application layer gateways for BITnet, MFEnet, and Internet (ESnet) Simple Mail Transfer Protocol (SMTP) electronic mail are provided at both Lawrence Berkeley Laboratory (LBL) and Argonne National Laboratory (ANL). In addition, LBL provides an electronic mail gateway from DEC VMS mail to the other electronic mail protocols mentioned above.

4.4 Network Operations Control Center

The ESnet Network Operations Control Center (NOCC) (Figure 4-9) provides 24-hour per day monitoring and control capabilities for the various network components that comprise ESnet. The Control Center is operated by the NERSC Engineering Group and the Supercomputer Operations staff. The Control Center staff continuously monitors the ESnet backbone routers to verify the network's integrity and to routinely gather statistics for troubleshooting and long term planning.

Electronic mailboxes for network information, network operations, and trouble calls are provided. An on-line trouble ticket system exists such that all reported problems are properly tracked.

4.5 Documentation

Specifications for ESnet IP and DECnet routing are available via anonymous file transfer protocol (FTP) from the ESnet Information Server (nic.es.net) in subdirectory [anonymous.specs].

National MFE Network 1991

NERSC

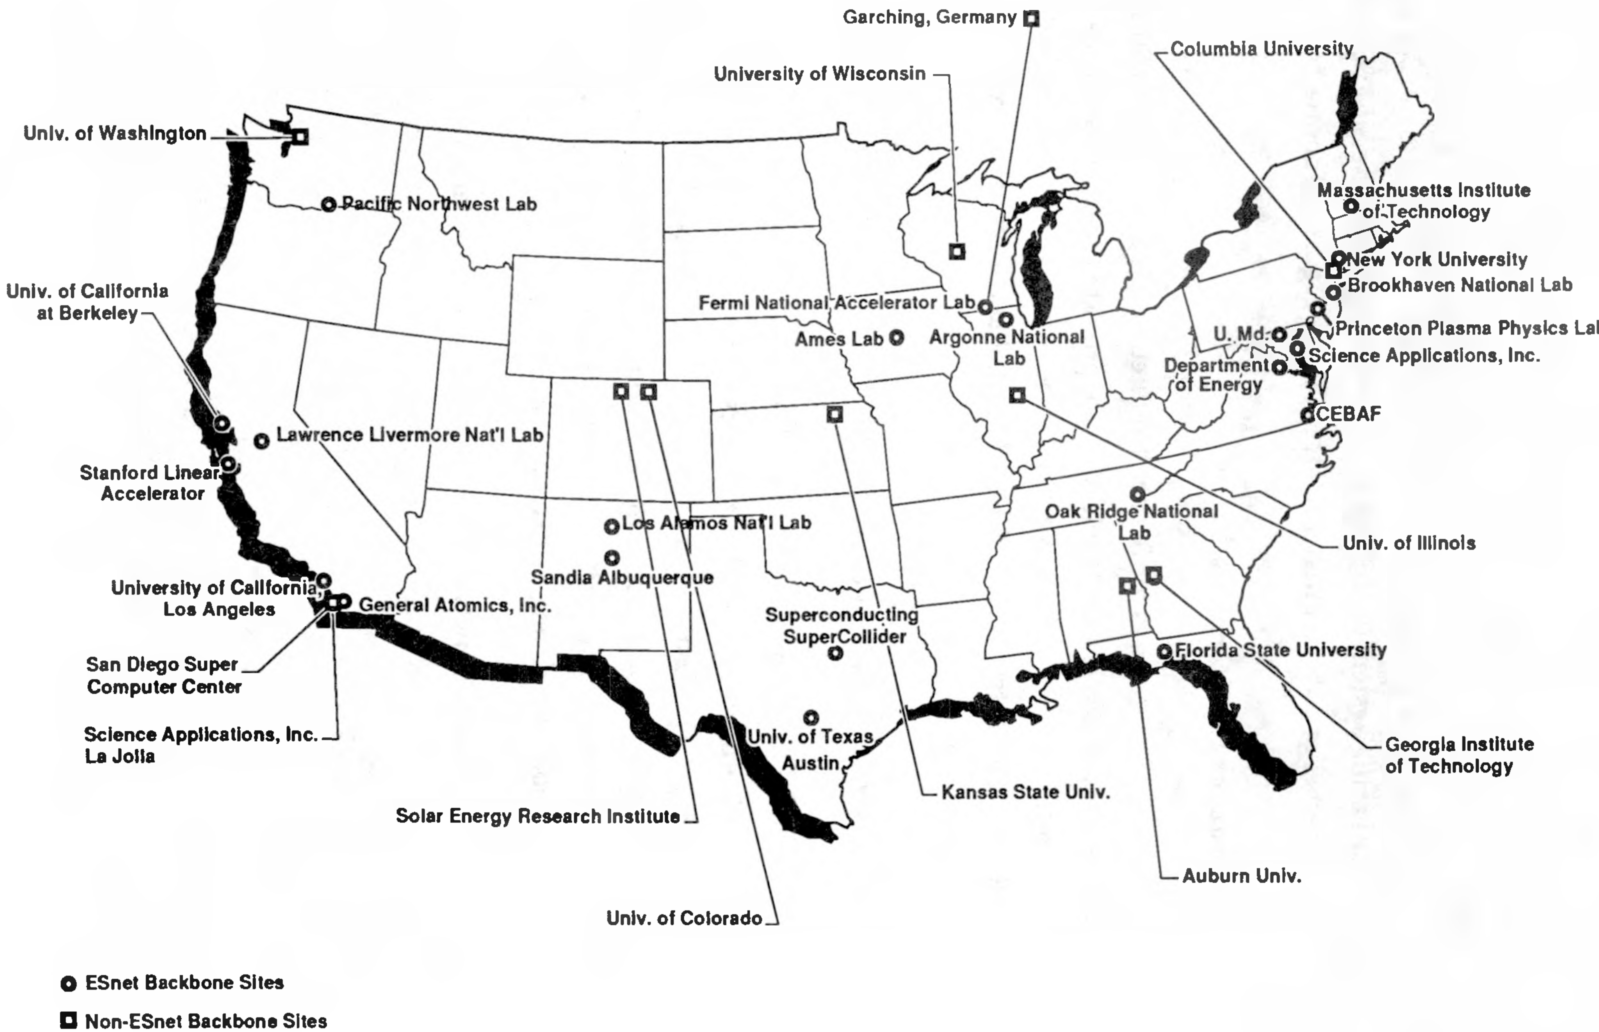

1/3/91

Figure 4-8. National MFE network, 1991.

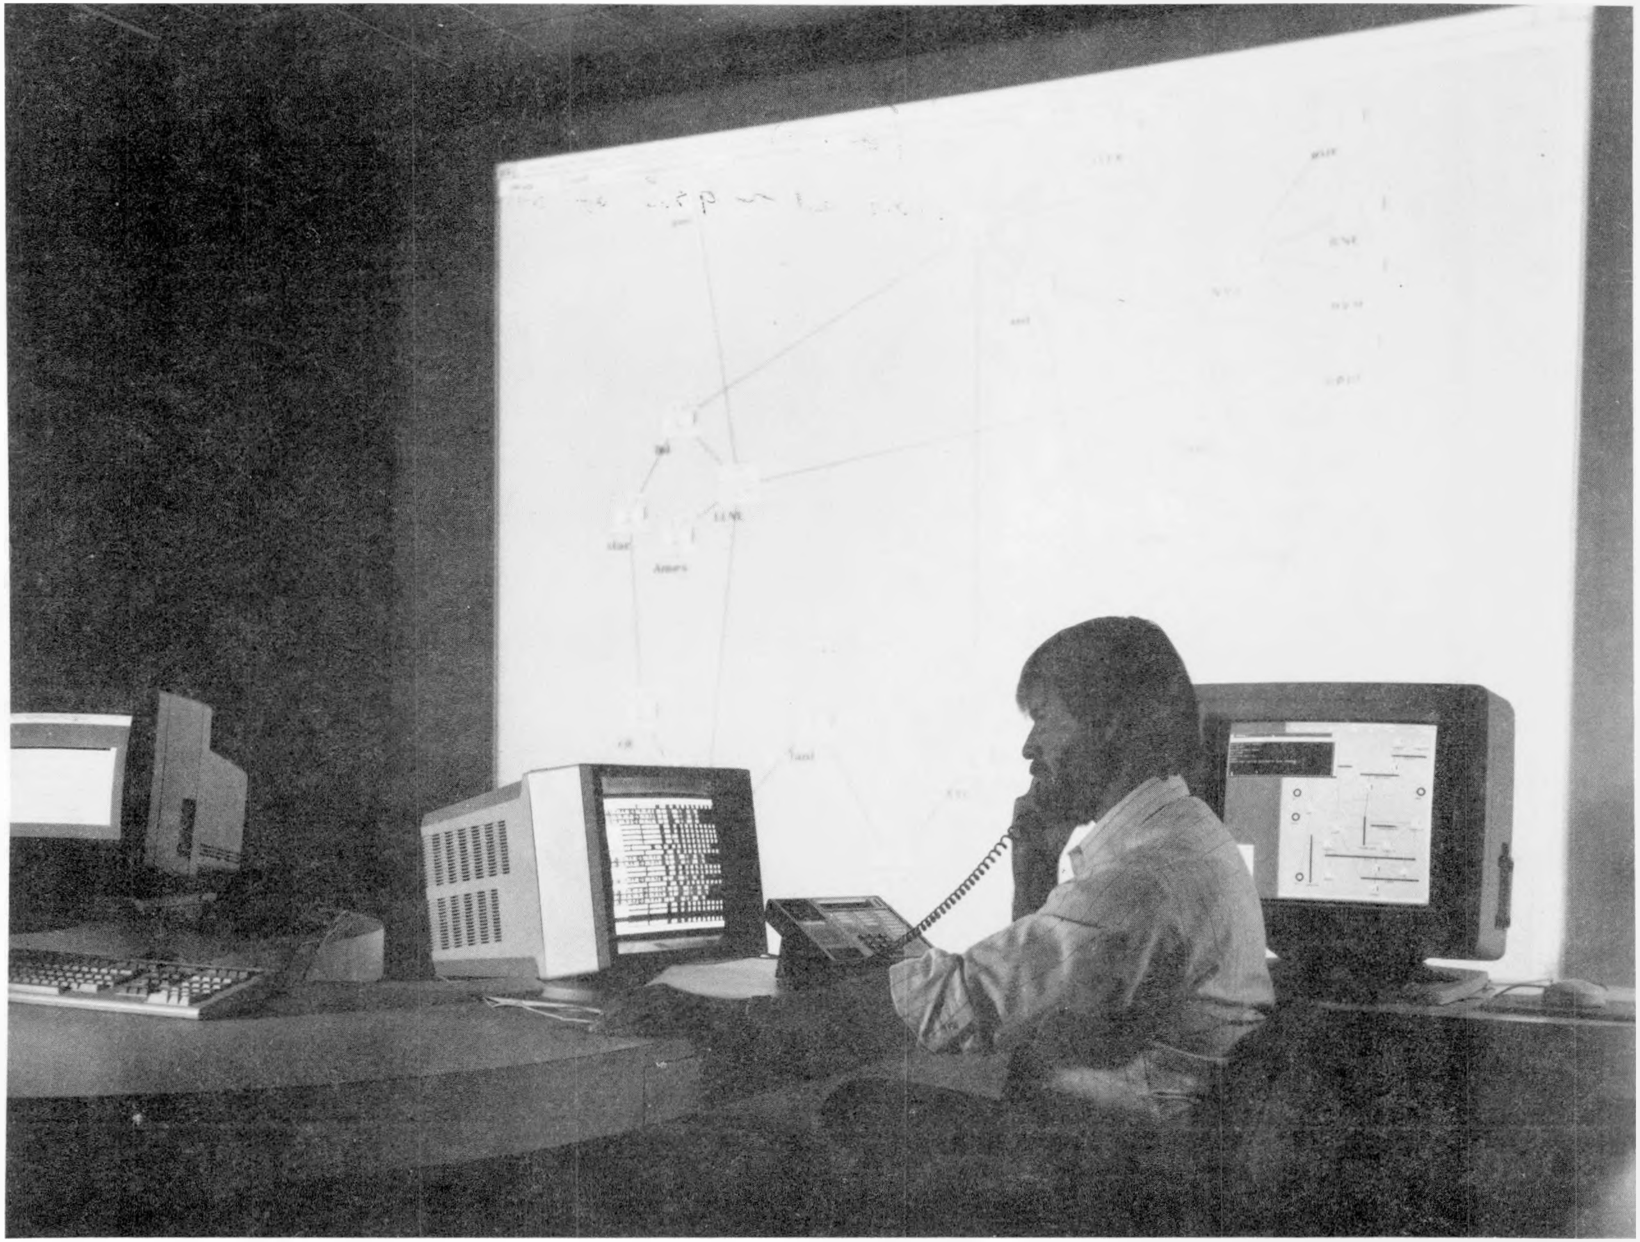

Figure 4-9. ESnet Network Operations Control Center.

The ESnet Policy Document, currently in draft form, will be made available following review and finalization.

Major Energy Research Programs and Requirements 5

Section 5: Major Energy Research Programs and Requirements

5.1 Program Introduction

The following subsections of the ESnet Program Plan contain information specific to the individual programs of the Office of Energy Research, and their related scientific computer networking needs. Information from the *FY 1990 Programs of the Office of Energy Research* [1] is extracted here as an introduction to the individual program subsections.

"The Office of Energy Research manages programs that provide basic science support for the several energy technologies and spin-offs to non-energy fields as well as advancing understanding in general science and training future scientists.

"Energy Research provides insights into fundamental science and associated phenomena and develops new or advanced concepts and techniques. Research of this type has been supported by the Office of Energy Research and its predecessors for over 40 years and includes research in the natural and physical sciences, including high energy and nuclear physics; magnetic fusion energy; biological and environmental research; and basic energy sciences research in the materials, chemical, and applied mathematical sciences, engineering and geosciences, and energy biosciences. These basic research programs help build the science and technology base that underpins energy development by Government and industry.

"Funding for these Energy Research programs totaled slightly over $2.1 billion in Fiscal Year 1989. These research programs support several thousand individual projects carried out by scientists from laboratories, universities, industry, and other research facilities throughout the United States. Most of these projects are carried out at the 9 major multiprogram and 16 single-purpose laboratories that conduct the Energy Research programs. The major research facilities at these laboratories have enabled scientists to continue to expand our technical knowledge base in the many areas of research pursued by the Office of Energy Research."

The relationship of the various program offices is shown in Figure 5-1.

Reference: 1. *FY 1990 Programs of the Office of Energy Research* (in press).

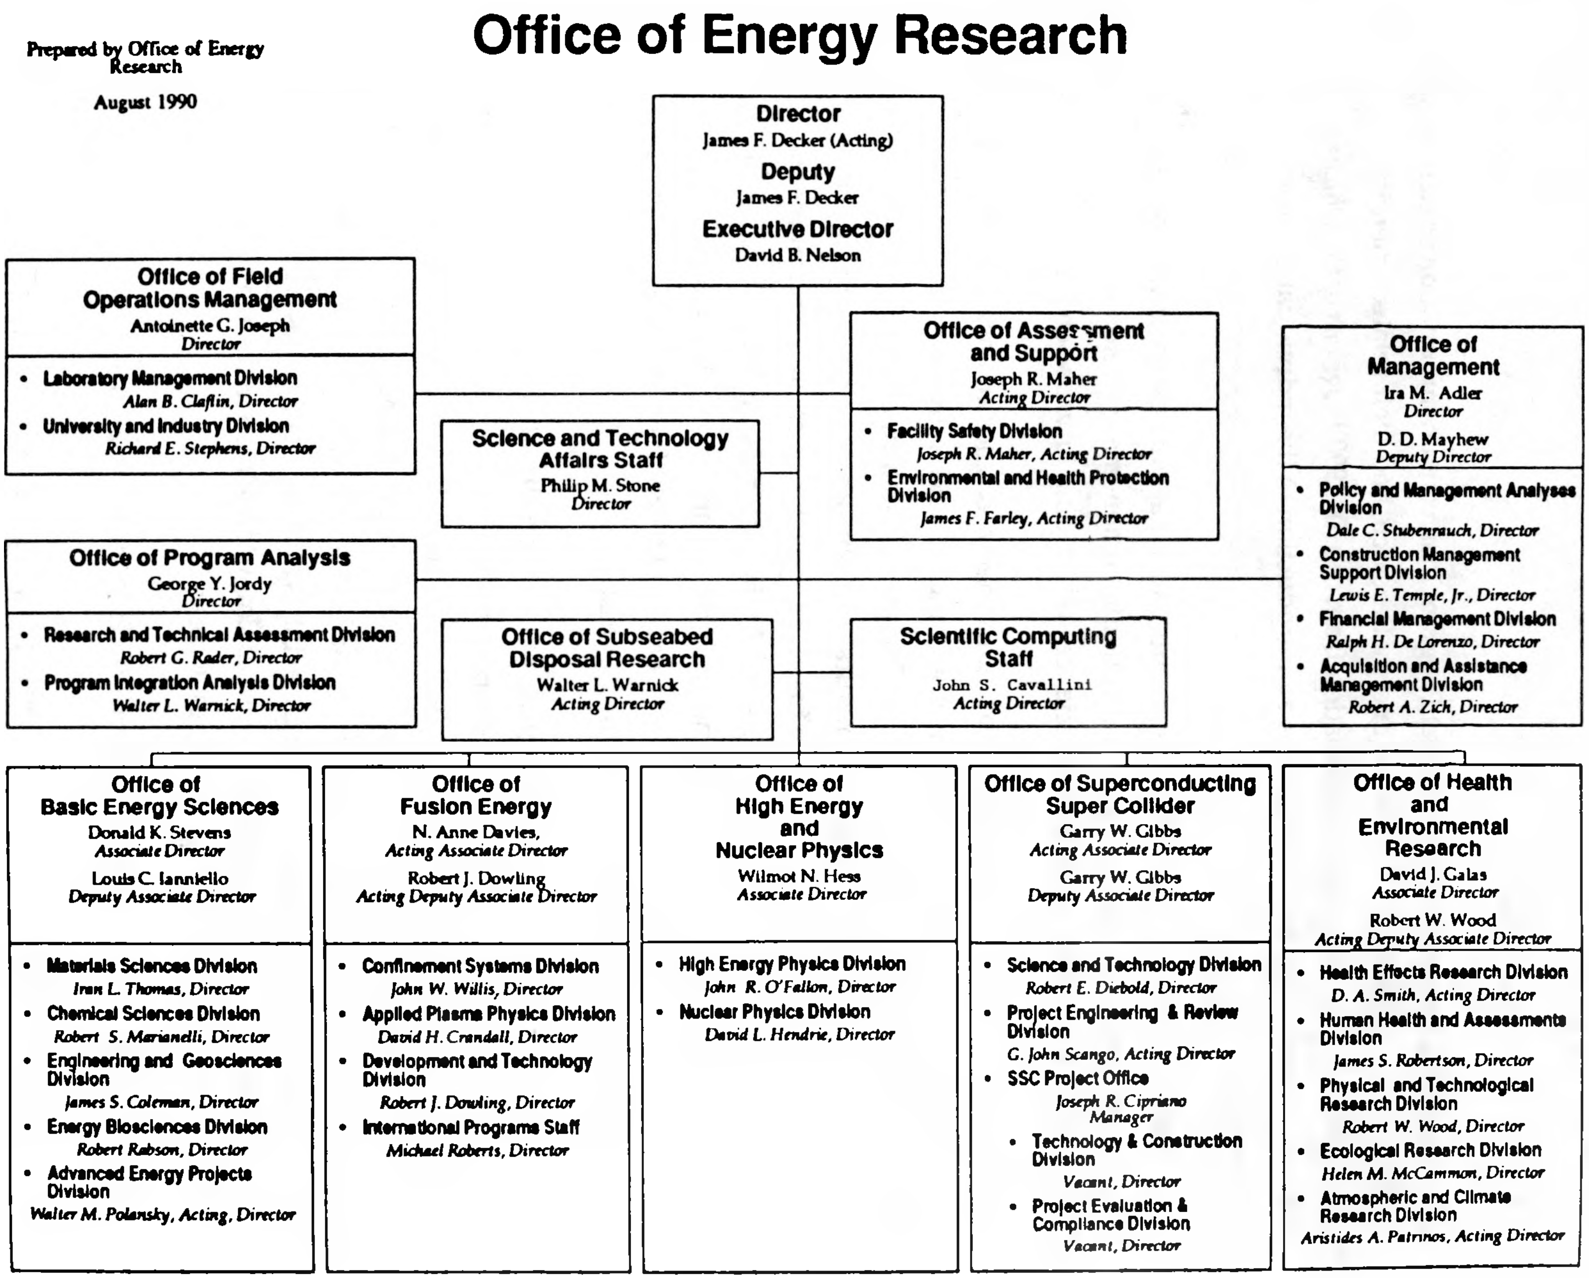

Figure 5-1. Organization chart, Office of Energy Research, Department of Energy.

5.2 Basic Energy Sciences

Materials research, which encompasses solid state physics and metallurgy and ceramics, is concerned primarily with the properties of solids. Materials Sciences research is traditionally characterized as small science with large numbers of individual investigators, but 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 Reference 1. Approximately two-thirds of the budget is spent through DOE laboratories, with the remainder spent at university-based programs, including both laboratories and individual grants. A list of user facilities appears in Table 5-1.

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 that 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 the current ESnet in this regard. 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.

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

Table 5-1. BES User Facilities.

Major User Facilities
In Use
National Synchrotron Light Source, Brookhaven National Laboratory
High Flux Beam Reactor, Brookhaven National Laboratory
High Flux Isotope Reactor, Oak Ridge National Laboratory
Intense Pulsed Neutron Source, Argonne National Laboratory
Los Alamos Neutron Scattering Center, Los Alamos National Laboratory
Stanford Synchrotron Radiation Laboratory, Stanford University
Under Construction
Advanced Light Source, Lawrence Berkeley Laboratory
Advanced Photon Source, Argonne National Laboratory
Other User Facilities
National Center For Small Angle Scattering Research, Oak Ridge National Laboratory
Electron Microscopy Center for Materials Research, Argonne National Laboratory
Shared Research Equipment Program, Oak Ridge National Laboratory
Center for Microanalysis of Materials, University of Illinois
Surface Modification and Characterization, Oak Ridge National Laboratory

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 for advances from 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 databases 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 to software so that preliminary data can be analyzed quickly and experiments altered in progress.

A second example of future uses of networking is access to 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 database 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 Sciences 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 experimentalists is expected in view of the importance of facilities at remote sites.

Networks and network services are also an important enabling technology for multiple-site collaborations among researchers. The use of electronic mail, file transfer, and access to remote databases has greatly enhanced productivity in BES research programs. These collaborations will become more practical and more productive when network-based services (such as distributed analysis and visualization, distributed file systems, video conferencing, and multiuser document processing) becomes available. Exploration and implementation of these emerging technologies are important parts of a comprehensive network development plan.

Reference:

1. *Materials Sciences Programs, FY 1989,* DOE/ER-0447P.

5.3 High Energy Physics

5.3.1 Introduction to High Energy Physics

High Energy Physics (HEP) as a discipline is concerned with the structure of matter and forces at the most fundamental level. The quest for understanding smaller and more basic components of matter has paradoxically required even higher-energy particles as its instruments (see Figure 5-2). Apart from an energetic theoretical component, most HEP investigations are carried out at the major accelerator centers, which in the U.S. are the Fermi National Accelerator Laboratory (FNAL), the Stanford Linear Accelerator Center (SLAC), Brookhaven National Laboratory (BNL), and Cornell University's Wilson Synchrotron. The Superconducting Super Collider Laboratory (SSCL) is just beginning to build the next major U.S. accelerator, which is anticipated to turn on in 1998. U.S. physicists are active users of accelerators abroad as well, principally the European Organization for Nuclear Research (CERN) in Geneva, Switzerland; the German Electron Synchrotron Laboratory (DESY) in Hamburg, Germany; the National Laboratory for High Energy Physics (KEK) in Tsukuba, Japan; and the Serpukhov Laboratory in Serpukhov, USSR.

Experiments at these accelerator centers are major enterprises, typically involving 100 to 500 physicists (and at least as many engineers and technicians during the building stages) and requiring five to fifteen years from conception to final data taking. The physicists achieve these numbers by collaborations between several institutions, typically with an international mix. Collaborations form primarily because of shared interests in a certain approach to physics with little concern for geographic proximity. As a result, all forms of communications, particularly computer networking, are crucial to the ability of collaborations to function smoothly (or at all). Computer networking is of particular importance because HEP experiments are of a complexity that can only be handled by computers at essentially every stage of the operation. Large codes are written to acquire, store, and analyze large samples of data. Each of these operations will typically involve collaborators at widely separated institutions. Fast, reliable, sophisticated networks are required to make the joint effort possible.

The large accelerator centers necessarily provide some centralization and focus to HEP networking needs, but another (smaller) component of HEP research works without accelerators, either using cosmic rays (and the cosmic accelerators that produce them) or

HIGH ENERGY PHYSICS

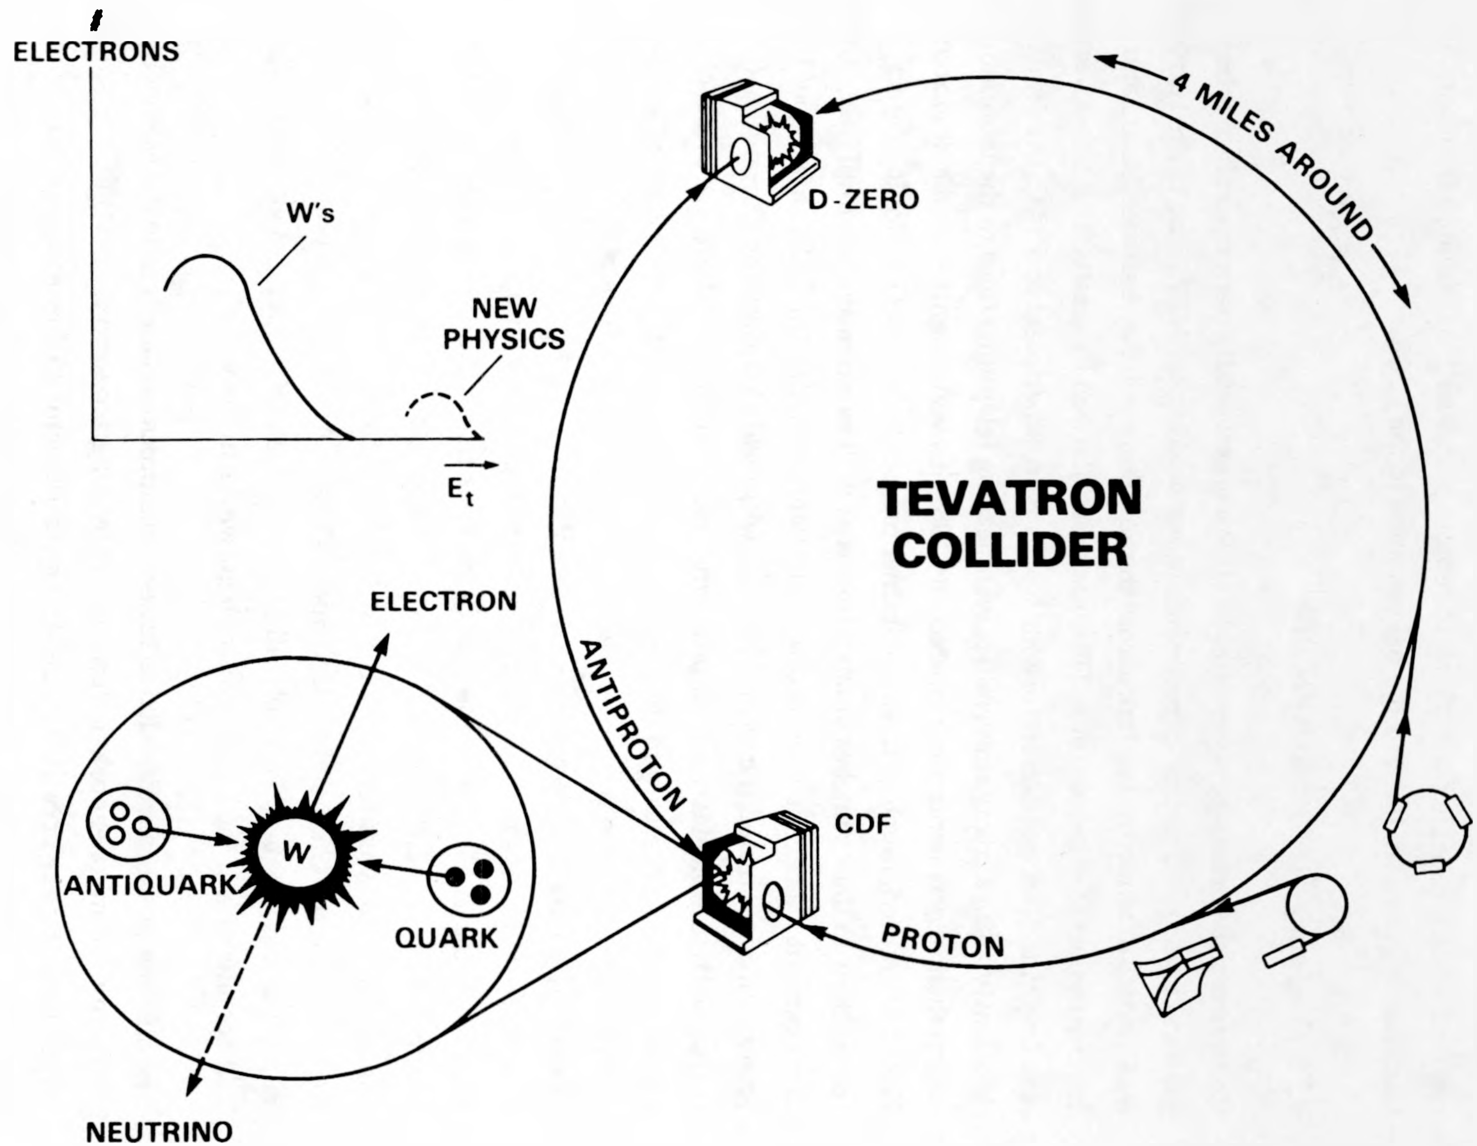

Figure 5-2. A HEP Tevatron collider.

looking for rare spontaneous events that require the shielding of a mountain or a mine in order to be sensitive enough to a tiny signal. This part of the work produces a requirement for reliable networking to remote places such as Sudan, Minnesota, the Gran Sasso Tunnel in Italy, or the Baksan mine in the USSR.

5.3.2 *Computing and Networking in HEP*

Computer networking between sites in HEP has essentially been used since it became practical to do so. The first experiments in networking for HEP came in the late 1970s, making use of a microwave link across the San Francisco Bay between SLAC and the Lawrence Berkeley Laboratory. This was then followed by a satellite link between SLAC and Argonne National Laboratory for work on an experiment at SLAC. At the same time (early 1980s), many universities were leasing telephone lines to the laboratories at which their experiments were located. By 1985, it was recognized that a *de facto* network had been created, much of it using DECnet protocols, but also employing direct connections to data switches at laboratories as well as extensive use of BITnet, and that management was needed for its orderly use and growth. The DOE Division of High Energy Physics created the *ad hoc* HEPnet Technical Coordinating Committee (HTCC) to deal with technical issues and problems created by the rapid growth of a tool that had already become necessary to the successful execution of HEP experiments. In the same year, a subpanel on computing of the DOE's High Energy Physics Advisory Panel, chaired by Joe Ballam, recommended formal creation of a High Energy Physics network (HEPnet), with dedicated backbone lines running at a minimum of 56 Kbps replacing for long hauls the *ad hoc* interconnection of the lines that had been installed for use on particular experiments. While this proposal was considered, the *de facto* HEPnet continued to grow under the increasingly expert management of the HTCC until, together with the comparably sized Space Physics Analysis Network (NASA) and the European and Japanese HEPnets, all of which necessarily operated as one network, it had become one of the largest research networks in the world.

The HEPnet proposal coincided in time with a proposal to upgrade the MFEnet, which had grown from a network dedicated to providing access to the NMFECC supercomputers at LLNL into a general purpose network for the use of magnetic fusion and other ER researchers. The Director of Energy Research suggested combining the two initiatives for efficiency and greater functionality into what has become ESnet. In order to evaluate this approach and more generally to map the future of networking in

HEP, the HEPnet Review Committee (HRC) was formed in 1987, charged with evaluating the needs of HEP networking and making recommendations on future directions, including the relationship with ESnet. The 1988 HRC report [1] recommended making use of the combined ESnet initiative but cautioned that a combined enterprise would necessarily have different priorities from those of HEP alone. The HRC therefore recommended an ongoing management structure for networking within HEP and noted that HEP might find it necessary to provide its own supplemental lines when and if ESnet's broader mandate did not allow it to supply connections or bandwidth that HEP found absolutely necessary. The HRC report found that T1 bandwidths on several backbone connections would be needed for HEP use alone by 1989, and several times that speed by 1992.

As ESnet has brought its T1-based network on line in 1990, most of the concerns raised in the HRC report have been addressed. The T1 lines have provided adequate bandwidth for immediate needs, and have even allowed the testing of video conferencing between LBL, FNAL, and the SSCL. The use of cisco routers on the network has provided the necessary flexibility in the implementation of protocols; namely, both internet protocol (IP) and DECnet are being transported, and residual X.25 is available where needed. Finally, there has been a high level of cooperation between the ESnet management and the new HEPnet Manager's office at FNAL.

5.3.3 HEP Networking Requirements in the 1990s

Services. The basic services of computer networking have been integrated into the functioning of HEP experiments and planning for the next generation of experiments, as well as phases of theoretical work. Electronic mail, file transfer, virtual terminal service, remote access to files, and remote job submission and monitoring tools must operate with complete reliability and at high speeds. These services are needed universally to interface with DECnet protocols and for a subset over BITnet. Increasingly, transmission control protocol (TCP)/IP is also being used by HEP, and these services are needed to interface with TCP/IP as well.

Some new services are needed urgently. A complete, readily accessible directory of institutions and individuals is a service whose lack significantly reduces the utility of the present network. Neither ESnet nor HEPnet management has been able to provide this service yet, in spite of recognition of the need for some time. It should become an urgent priority. Recent trials of video conferencing have made it clear that this service will be needed by HEP and that its use will expand virtually as fast as network bandwidth permits. It seems likely that within two years, as much bandwidth will be devoted to video as to data.

Use of graphical windowing access to remote computers (X11, Motif, ...) is in the experimental stages now but can be expected to grow rapidly within HEP. Bandwidth requirements for interactive use will grow by an order of magnitude within 2–3 years for this reason alone.

Connectivity. HEP groups at laboratories and universities need access to experiments and databases at laboratories and at only slightly reduced bandwidth access to each other. The field and collaborations within the field are highly international, and connections to major international locations such as Europe and Japan need to be as good as domestic connections. To the extent that this is not true, the efficiency and productivity of present collaborations is demonstrably reduced. As an indication of connections needed, Table 5-2 (taken from the HRC Report [1]) lists the major HEP collaborations participated in by U.S. institutions in 1987, listed by the site of the experiment. This list is still representative of the distribution of HEP physicists, but it is now augmented by the addition of the SSC program. Approximately 1000 U.S. and 1000 foreign physicists have signed the initial Letters of Intent in doing experiments at the SSCL.

Even though the SSCL will not be completed before 1999, collaborations to build the necessary large detectors have already been formed. Fully capable networking will be needed among the collaborating laboratories and universities, starting immediately as collaborators design their detectors and make plans to begin construction. SSCL collaborations will be even more inclusively international than present-day HEP collaborations. High-performance networking is needed not only to Western Europe and Japan but also to the Soviet Union and China. Basic links will be needed to many other regions of the world. To the extent that administrative barriers still exist to networking to some countries, it should be a strong goal of ESnet to remove the restrictions.

Performance. The HRC report [1] was able to calculate required bandwidths that have approximately matched subsequent experience by counting the users needing to access labs on each link via compilations. Current measurements indicate that the peak and average usage of ESnet backbone lines by HEP is already heavy and is expected to reach saturation on some links soon. The HRC bandwidth estimates can be extrapolated by using the updated assumptions (estimated for 1992) that each interactive user requires 50 Kbps during the 10% of the working day that the user is interactively using the Laboratory's computing facility and that each user transfers 20 Mbytes per working day to and from the Laboratory. The numbers are then doubled to account for all traffic between the users rather than just that between users and laboratory. The results of these assumptions are as follows:

- Links to Europe and Japan must reach 1 Mbps by 1992.
- Tail circuits to individual sites must be at least 56 Kbps and in some cases 1 Mbps by 1992.
- Links to the Soviet Union and China are needed now at 64 Kbps and should grow to 500 Kbps by 1992.
- Backbone bandwidths within the U.S. must be several Mbps starting in 1992 and must reach T3 speed (45 Mbps) by 1994.

Note that these bandwidths are calculated for HEP use alone. Total required bandwidth must be calculated by summing all of these program descriptions. Although detailed calculations are not possible now, all requirements calculated for 1992 will increase by at least an order of magnitude and probably more by the end of the decade.

Planning and Management. It has been emphasized above that networking is now a required support capability for all stages of both experimental and theoretical work in HEP. Not only must the required bandwidth and services be provided promptly, but also management must be sufficiently responsive to prevent interruptions of service. Further, management must provide forecasting of performance and connectivity requirements adequate to provide what is needed before its lack hampers the scientific program. The required growth in services is fueled both by quantitative advances in computing and by qualitatively new capabilities that are made possible, as well as by expansion of the program.

Reference:

1. HEPnet Review Committee, *High Energy Physics Computer Networking,* DOE/ER-0372, June 1988.

Table 5-2. Number of remote HEP collaborators who could have used a HEP backbone to their experimental site in 1987.

Expt -	Institute	# people	KEK-LBL	LBL/SLAC-FNAL	FNAL-CERN	FNAL-DESY	FNAL-Italy	FNAL-BNL/MIT
CDF -	Brandeis	8						8 >
(FNAL)	Frascati	8					8 >	
	Harvard	11						11 >
	KEK	4	< 4	< 4				
	LBL	8		< 8				
	Penn U	9						9 >
	Pisa	20					20 >	
	Rockefeller	7					7 >	
	Rutgers	6						6 >
	Tsukuba	21	< 21	< 21				
SLD -	Bologna	1		1 >			1 >	
(SLAC)	Boston U	3		3 >				3 >
	Brunel	5		5 >	5 >			
	Cincinnati	4		4 >				
	Nevis	10		10 >				10 >
	Ferrara	2		2 >			2 >	
	Frascati	5		5 >			5 >	
	Illinois	4		4 >				
	MIT	18		18 >				18 >
	U. Mass	2		2 >				2 >
	Padua	5		5 >			5 >	
	Perugia	10		10 >			10 >	
	Pisa	3		3 >			3 >	

Expt -	Institute	# people	KEK-LBL	LBL/SLAC-FNAL	FNAL-CERN	FNAL-DESY	FNAL-Italy	FNAL-BNL/MIT
MKII -	Indiana	8		8 >				
(SLAC)	J.Hopkins	7		7 >				7 >
	Michigan	8		8 >				
D0 -	BNL	9						9 >
(FNAL)	Brown	4						4 >
	Columbia	3						3 >
	LBL	9		< 9				
	Maryland	2						2 >
	NYU	13						13 >
	Penn U	5						5 >
	Saclay	7			7 >			
	SUNY	28						28 >
	Yale	2						2 >
AMY -	Rochester	16	16 >	16 >				
(KEK)	LSU	5	5 >	5 >				
	Ohio S.	5	5 >	5 >				
	Rutgers	10	10 >	10 >				10 >
	S. Carolina	6	6 >	6 >				
	Virginia T	6	6 >	6 >				
	UCDavis	6	6 >					
ALEPH -	FSU	3			< 3			
(CERN)	Wisconsin	26			< 26			
L3 -	CalTech	12		< 12	< 12			
(CERN)	CarnegMel	8			< 8			8 >
	Harvard	5			< 5			5 >
	Hawaii	1		< 1	< 1			
	J Hopkins	4			< 4			4 >
	Mich U.	8			< 8			8 >
	MIT	31			< 31			31 >
	N Eastern	7			< 7			7 >
	Ohio S.	8			< 8			8 >
	Okla U.	1			< 1			1 >
	Princeton	11			< 11			11>

Table 5-2 (continued)

Expt -	Institute	# people	KEK-LBL	LBL/SLAC-FNAL	FNAL-CERN	FNAL-DESY	FNAL-Italy	FNAL-BNL/MIT
OPAL - (CERN)	Chicago	8			< 8			
	Maryland	12			< 12			
	Montreal	5			< 5			
	RC Canada	8			< 8			
	UCRiverside	5		< 5	< 5			
ZEUS - (DESY)	Carleton	1				< 1		
	McGill	3				< 3		
	Toronto	8				< 8		
	York (Can)	2				< 2		
	Ohio S.	3				< 3		
	ANL	5				< 5		
	Nevis	7				< 7		7 >
	Penn S.	3				< 3		3 >
	Virg Polyt	2				< 2		2 >
	Wisconsin	3				< 3		
	Illinois	2				< 2		
H1 - (DESY)	UCDavis	10		< 10		< 10		
	Houston	4				< 4		
	Northeastern	3				< 3		3 >
HELIOS - (CERN)	BNL	6			< 6			6 >
	LANL	6			< 6			
	McGill	5			< 5			
	Montreal U	7			< 7			
	Pitt U	6			< 6			6 >
	KEK	1	< 1	< 1	< 1			
PS-183 - (CERN)	UCIrvine	7		< 7	< 7			
	Penn S.	8		< 8				8 >
	New Mex	6		< 6				
PS-200 - (CERN)	LANL	18			< 18			
	Rice	1			< 1			
	Texas A&M	5			< 5			
	Kent St	1			< 1			
	Case West	1			< 1			
	AMES	1		< 1	< 1			
UA-1 (CERN)	Harvard	7			< 7			7 >
	MIT	5			< 5			5 >
	UCRiverside	11		< 11	< 11			
	Victoria	6		< 6	< 6			
	Wisconsin	8			< 8			
-704- (FNAL)	Kyoto	11	< 11	< 11				
	Saclay	6			6 >			
	Trieste	10					10 >	
	KEK	2	< 2	< 2				
-632- (FNAL)	Eng+CERN	43			43 >			
	Munich	4				4 >		
	Tufts	6						6 >
	LBL	4		< 4				
	Hawaii	6		< 6				
-636- (FNAL)	MIT	11						11 >
	Brown	4						4 >
	Tohoku	12	< 12	< 12				
-652-	Columbia	11						11 >
-653- (FNAL)	UCDavis	6		< 6				
	KOBE	8	< 8	< 8				
	Nagoya	12	< 12	< 12				
	Osaka	5	< 5	< 5				
	Other Japan	10	< 10	< 10				

Table 5-2 (continued)

Expt -	Institute	# people	KEK-LBL	LBL/SLAC-FNAL	FNAL-CERN	FNAL-DESY	FNAL-Italy	FNAL-BNL/MIT
-665-	UCSD	5		< 5				
(FNAL)	Freiburg	5				5 >		
	Harvard	6						6 >
	Maryland	8						8 >
	MIT	6						6 >
	Munich	10				10 >		
	Wash U	8		< 8				
	Yale	4						4 >
-687-	Milan	11					11 >	
(FNAL)	Frascati	6					6 >	
-690-	Nevis	6						6 >
(FNAL)	Mass U	4						4 >
-769-	Northeastern	2					2 >	
(FNAL)	Tufts	5					5 >	
	Yale	1					1 >	
-733-	MIT	12						12 >
BNL838	Minn	5						< 5
(SIN)	all USA	19		< 3	< 19			9 >
	BC Canada	6			< 6			
MACRO	Cal Tech	4		< 4	< 4			
(Gran	Indiana	3			< 3			
Sasso)	Northeastern	9			< 9			9 >
	Michigan U.	5			< 5			
	Texas A&M	2			< 2			
	Virginia T	3			< 3			

GRAND TOTALS

	KEK-LBL	LBL/SLAC-FNAL	FNAL-CERN	FNAL-DESY	FNAL-Italy	FNAL-BNL/MIT
WESTWARD TOTAL	< 86	< 192	< 330	< 56	< 0	< 5
EASTWARD TOTAL	54 >	160 >	64 >	19 >	89 >	384 >
TOTAL PEOPLE (EAST+WEST)	140	352	394	75	89	389

5.4 Magnetic Fusion Energy

5.4.1 Introduction

The long term goal for the Magnetic Fusion Energy (MFE) program is the development of fusion energy as a practical energy source. The approach is to use strong magnetic fields to confine plasmas of sufficient temperature and density for the fusion reactions to proceed. To reach this goal, it has been necessary for the program to develop the science of plasma physics, a field closely related to both fluid and statistical mechanics. The turbulent nature of fusion-grade plasmas and the complicated geometry of fusion experiments make closed-form theoretical treatment of the problem impossible. Fusion experiments require the acquisition and analysis of very large numbers of plasma parameters as functions of space and time, and of energy and time, as well as the acquisition and analysis of video images. As a result, both experiments and theory have become computationally intensive.

In this country, work on MFE is spread over almost 100 sites, with concentrations at about a dozen national laboratories, universities, and private contractors. Research on a similar scale is carried out in Europe, Japan, and the USSR. Progress is linked to a handful of large experiments, at facilities too expensive to be duplicated, such as the Tokamak Fusion Test Reactor (TFTR) at the Princeton Plasma Physics Laboratory (PPPL). As the size of the experiments has increased, the number and importance of collaborations between small groups, mostly located at universities, and the large groups at the experimental sites has also increased. Future experiments will be the product of collaborations on the national (CIT, the Compact Ignition Tokamak, at PPPL [1]) and international (ITER, the International Thermonuclear Experimental Reactor [2]) scale. Collaborations for these experiments will be spread over the entire MFE community. Designs for these experiments are already under way by national and international teams, respectively.

5.4.2 Computing and Networking in MFE

The computing environment for most fusion researchers was established with the creation of the National Magnetic Fusion Energy Computer Center (NMFECC) in 1974, recently renamed the National Energy Research Supercomputer Center (NERSC) to reflect its expansion to a broader range of energy research. Starting with a CDC 7600, the Center expanded its facilities by acquiring a variety of Cray supercomputers. The supercomputers presently at the Center are a Cray-XMP and three Cray-2's. High speed communications among major research centers and NERSC have until recently been provided by MFEnet, consisting largely of 56-Kbps satellite links. MFEnet is currently being upgraded to consist mainly of 1.5-Mbps T1 fiber optic land lines as part of ESnet. Smaller sites are still connected to the nearest major site through slower land lines or even simple dial-up lines.

Problems requiring supercomputers for their solution occur throughout the MFE program. Large scale engineering codes are used in many areas of analysis of magnetic confinement devices: thermal, structural, neutronic, electromagnetic, power, and overall system design. Transport codes are composed of many modules containing a good deal of physics and can result in lengthy calculations. Three-dimensional equilibrium computations are needed for stellarators and for toroidal ripple in a tokamak. To model the wide range of instabilities in a fusion device, three-dimensional fluid and kinetic calculations with sufficiently fine resolution are required. At present, one of the great "Grand Challenge" problems in the MFE program is the understanding of plasma turbulence and heat transport. The design of new machines and of proposed reactors depends critically on the predictions of the magnitude of this transport. To date, the fusion community has relied mainly on empirical scaling laws to predict transport. More accurate predictions are a widely recognized goal of the community.

High speed communications are an essential part of the solution to these many computational problems. Computing power at NERSC and the scope of the calculations performed there have increased substantially over the last sixteen years. One single state-of-the-art turbulence calculation can now generate several gigabytes of data describing the evolution of the various physical quantities of interest in space and time. Traditionally, these data were archived and analyzed at NERSC, and a minimal amount of data file transfer to, and printing and plotting at, the local site was performed. The advent of affordable local mass storage devices, such as optical jukeboxes, and powerful desktop workstations for efficient post-processing, in-depth analysis, and graphical display of data, will mean an increase of these activities at the local site and also an increase in the data communications requirements between the local site and NERSC. Even greater pressures on network bandwidth and availability are expected with the move towards a degree of "seamless" integration between the user's local system and the resources at NERSC.

Large fusion experiments produce in excess of 200 gigabytes of raw data per year. For machines like CIT or ITER, the number would likely be an order of magnitude or more larger. The task of acquiring, archiving, displaying, and analyzing these data requires substantial manpower and hardware. Computer staffs at the sites struggle to prevent these tasks from limiting the rate at which research is carried forward. Traditionally, virtually all of these data were archived and analyzed locally. In recent years, collaborations required much more of these data to be either transferred or processed by remote users. The creation and expansion of the National Tokamak Database at UT Austin and implementation of the CIT and ITER R&D programs will also increase the data communications requirements for the MFE program.

During the design phase for the next-generation fusion devices, CIT and ITER, large-scale engineering analysis will take a significant share of the computing resources. This analysis includes electromagnetics, structural, neutronics, and heat transfer calculations. Since the design groups are widely distributed nationally and internationally, communication of models and results of simulations and analyses will be an important requirement for this network. Planning is under way for the ITER EDA (Engineering Design Activity), a five-year program to produce a fully engineered design. The site for the EDA headquarters, which will house about 300 professional staff, will be chosen within the next six months. These staff members will have to coordinate the activities of many hundreds of co-workers scattered among the collaborating ITER nations. Regardless of the choice of site, connection to high-performance international networks will be crucial for the success of this mission.

By the time CIT is in operation, users at their home institutions will require transparent access to the entire archive. The CIT physics design team has proposed a system of "distributed control rooms" from which remote groups could control the characteristics of diagnostic equipment, and monitor and analyze data as they are acquired. This system obviously requires that such groups have excellent communications with the local computer system, physics group, and machine operations staff, as well as with other physics groups at remote sites.

5.4.3 Networking Requirements: Connectivity

Overall approximately 100 sites in the U.S., 65 sites in Europe, 45 sites in Japan, and several sites in South America, China, and Australia are involved in MFE research. The principal sites are shown in Table 5-3.

The highest priority for network connectivity in the U.S. MFE program has been access to the supercomputer center at LLNL. While this will likely be the case for the foreseeable future, requirements for connectivity between laboratories in support of collaborations and access to remote databases will continue to grow. Access to the site of the CIT device (which is proposed to come on line in 1997) will become more and more important as the decade proceeds.

Specific connectivity requirements are:

- Support access to NERSC at LLNL.
- Support experimentalists at remote sites operating diagnostics and performing experiments on TFTR and PBX at PPPL, ATF at ORNL, DIIID at GA, C-Mod at MIT, TEXT at UTA, JET at Culham, ASDEX and ASDEX-U at IPP Garching, TORE SUPRA at Cadarache, TEXTOR at Julich, FTU at Frascati, RFX at Padua, and JT-60 at Tokai.
- Support access to the tokamak database at UTA.
- Support communications among the CIT design team located at PPPL, ORNL, MIT, and GA.
- Support CIT physics and engineering R&D. [3]
- Support CIT diagnostics development.
- Support ITER design team at Garching.
- Support ITER R&D program. [4]
- Support gateways into Internet, SPAN, BITnet, etc.

Table 5-3. Principal sites involved in MFE research.

USA

Argonne National Laboratory
Columbia University
* General Atomics
Lawrence Berkeley Laboratory
* Lawrence Livermore National Laboratory
* Los Alamos National Laboratory
* Massachusetts Institute of Technology—Plasma Fusion Center
National Energy Research Supercomputer Center (at LLNL)
Naval Research Laboratory
New York University
* Oak Ridge National Laboratory
* Princeton Plasma Physics Laboratory
Sandia National Laboratories
Science Applications, Inc.
University of California at Berkeley
University of California at Los Angeles
University of California at San Diego
University of Illinois
University of Maryland
* University of Texas
University of Washington
* University of Wisconsin

International

Austria: IAEA Headquarters, Vienna

*Canada: Varennes—Hydro Quebec CCFM

*France: Cadarache—CEN

Germany:
* Garching (near Munich)—IPP, Max Planck Institute
* Julich—KFA
Karlsrühe—KFK

Italy:
* Frascati—CREF (ENEA)
* Padua—IGI

Japan:
* Kyoto—Kyoto Univ., Uji
* Nagoya—IPP
* Naka—JAERI
* Tokai—JAERI

* Switzerland: Lausanne—CRPP, EPFL

* UK: Culham/Harwell—JET, Culham Laboratory

*Indicates the sites of major MFE experiments.

5.4.4 Functional Requirements

At a minimum, the MFE program requires the standard network services, including electronic mail, file transfer, remote logins, remote printing and plotting, and support for bulletin boards. These services must be available for the supercomputers at NERSC. It is preferable if the connections are made on terrestrial lines to eliminate the delays inherent in satellite communication and to expedite services such as full screen editing. Support for TCP/IP and DECnet is required along with BITnet for e-mail to otherwise inaccessible (mostly foreign) sites. Lastly, some sort of network directory service is needed, supporting electronic white and yellow pages.

It is expected that developments under way will result in a new paradigm for computing, where activities (computing, visualization, file management, etc.) for a given task are shared over a network between several machines. Already, windowing terminals and workstations with a variety of distributed network services are being used in large numbers by the fusion program. This new paradigm for computing not only requires greater network bandwidth (and very high availability), it also requires a new class of functionality to achieve what is often referred to as "seamless" computing. These functions would include remote job entry, remote procedure calls, network file service, distributed editors, distributed databases, and so forth. There will also be a need for remote or distributed code management. The nature of the MFE program requires that these services be available on wide area as well as local area networks.

The multi-institutional and multinational nature of current design activities, and the increasing importance of collaborations in experiments and theory with the same broad participation, will require better tools for interpersonal communications. The CIT and ITER design teams are eager to hold video conferences between their widely scattered members and with other experts in the field. They are held back now only by the lack of available network bandwidth. Groups at the University of Wisconsin, Columbia, and MIT collaborating on experiments at the TFTR have requested funding to support, in essence, multimedia mail with voice and graphics. Such capabilities will need to be made more generally available.

It has been proposed that, as part of the CIT program, distributed control rooms be set up to support remote data acquisition and instrument control for groups located at their home institutions. (This would not of course extend to machine control.) The nature of the MFE experimental method requires close communication between all members of the physics teams acquiring and analyzing data for every experimental run. While it would be expected that many members of a national CIT physics team would be present on site for many runs, it is unlikely that they would all be present every day. Thus the proposal is to extend the control room out to the collaborating sites, linked with data, voice, and video.

5.4.5 Performance Requirements

The current ESnet, with its MFEnet and DECnet components, its T1 backbone, and its gateways to Internet and BITnet, are meeting most of the present needs of the MFE program. The most notable exception is the lack of video conferencing capabilities and the lack of high speed tail circuits to non-backbone sites. To meet present requirements, most tail circuits need to be upgraded, with ten going to full T1 service.

It is more difficult to extrapolate future requirements with any precision, since users will not even contemplate applications that cannot be supported. These requirements can be estimated by a number of methods. First, all the major sites have local area networks running at 10 Mbps and using (at least for peak loads) a fair fraction of this bandwidth. Many of these sites have plans to extend their local networks to 100 Mbps as fiber distribution data interface (FDDI) hardware becomes available. The services described above require local area network (LAN) functionality over the wide area network (WAN). Summing loosely over the sites and assuming that communications between sites accounts for approximately 10% of the local traffic, the WAN requirements can be estimated at 10–100 Mbps. A similar number can be reached by noting the proliferation of workstations on the desks of physics and engineering staffs. It is reasonable to assume that there will be on the order of 1,000 of these before very long. It has been estimated that when workstations are actively using network resources, they require about 100 Kbps. Allowing for a 10% utilization factor overall, with 10% of that on wide area networks, and multiplying by 10 for a peak to average load factor and 3–10 for increased cpu speedups over the next five years, one arrives at a figure of 30–100 Mbps. Finally, consider the requirements for multimedia e-mail, video conferencing, distributed computer-aided design (CAD), and so forth. It would seem that a conservative forecast would be for the equivalent of several tens of T1 lines by the end of five years.

References:

1. *Physics Aspects of the Compact Ignition Tokamak,* D. Post *et al.*, PPPL-2389.
2. *ITER Concept Definition,* International Atomic Energy Agency, 1989.
3. *CIT Physics R&D Requirements*, AE-900216-PPL-01.
4. *ITER Related Physics R&D Needs*, ITER-IL-PH-16-0-18.

5.5 Nuclear Physics

5.5.1 Programmatic Introduction

The central objectives of the Nuclear Physics program are understanding the interactions, properties, and structures of nuclei and nuclear matter and understanding the fundamental forces of nature as manifested in atomic nuclei. In Nuclear Physics, until about ten years ago, there were many small facilities available, generally one per institution, and many Nuclear Physics researchers performed experiments at their own machines. Most experiments in Nuclear Physics measured less than ten parameters per event and had 1–2 institutions per experiment. Nuclear Physics research tended to be localized, small, and independent. Networks were not essential.

In the last ten years Nuclear Physics has been moving to higher energies, and to more complex experiments at low energies, as the "easy" experiments pay off less and less. This has led to a similar situation to that which High Energy Physics has been facing. The number of facilities has been dropping and the size of collaborations is growing. Nuclear Physics now has experiments that measure thousands of parameters and have tens of institutions participating. Nuclear Physics experiments have matched or exceeded High Energy Physics experiments in several areas of complexity. The particle multiplicities in Nuclear Physics experiments at the CERN Super Proton Synchrotron (SPS) are higher than High Energy Physics expects to see at the Superconducting Super Collider (SSC). The Heavy Ion Spectrometer System (HISS) time-projection chamber at LBL has 16,000 pads, making it comparable to many High Energy Physics experiments.

The main differences between High Energy Physics and Nuclear Physics that affect the need for network resources are location, diversity, and scale. The upper end of diversity of individual requirements in Nuclear Physics matches well with that of High Energy Physics. Nuclear Physics does not overlap completely on experimental locations [the Continuous Electron Beam Accelerator Facility (CEBAF), for instance] (see Figure 5-3). Nuclear Physics' needs for everything other than connectivity and total bandwidth are the same.

SITE PLAN CEBAF

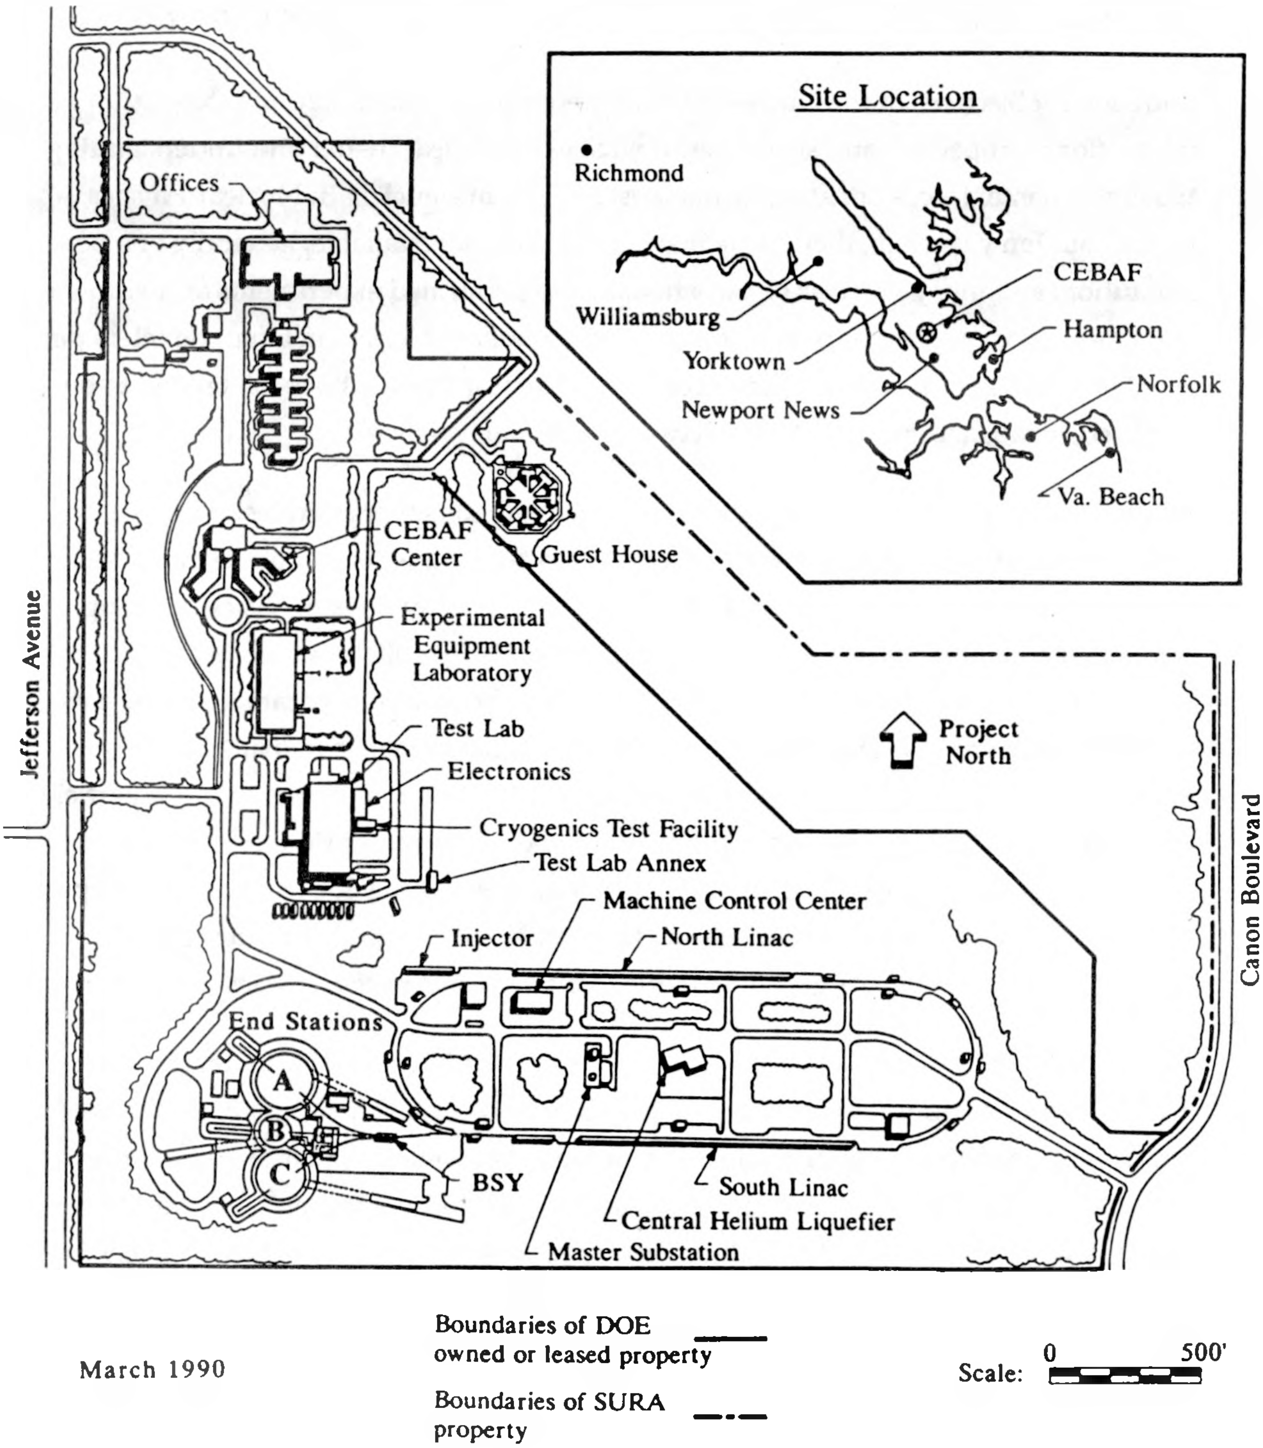

Figure 5-3. CEBAF site overview, 1990.

5.5.2 Programmatic Functional Requirements

Planning and construction of ESnet is an opportunity to provide "enabling technology." This means that as the needs of the Nuclear Physics research community grow, ESnet will have foreseen the needs and will have the required network in place, thus allowing rapid progress by the researchers. The Nuclear Physics Panel on Computer Networking (NPPCN) has considered the current and the enabling requirements of the Nuclear Physics community. In considering these needs, the panel has reviewed the "Report of the HEPnet Review Committee," DOE/ER-0372, [1] and finds that there is extensive overlap between the needs of the Nuclear Physics and the High Energy Physics communities. Both sets of requirements are discussed below.

Current Requirements

The following list contains network properties that are considered necessary by the Panel to serve the current needs of the Nuclear Physics community. Most of these properties are available in the current HEPnet and certainly form a basic requirement for ESnet in its early phase.

1. *Electronic Mail.* Mail exchange of short notes between all facilities and users on all possible networks should be quick (less than 1 minute). Both immediate and store-and-forward options should be available.

2. *File Transfer.* Files, either binary or text, need to be transparently (i.e., independent of internal representation) sent from node to node. Transmission rates of 100 Kbps are required now for existing Nuclear Physics experiments.

3. *Network Virtual Terminal.* Interactive response time should be less than 0.2 second for connections between two network nodes. Ideally, standard computer facilities should be available when using network virtual terminal access to a network node, including logging on and use of printers and graphics devices.

4. *Protocols.* Continued support of full "native mode" DECnet communication between DEC hosts is required to prevent severe budgetary and functional losses due to the current high Nuclear Physics usage of DEC software and hardware.

5. *Task-to-Task Communication.* Task-to-task communications, regardless of the type of host, are necessary. Examples that are currently available are batch job submission and file-directory searches.

6. *Connectivity.* Connections are required to all laboratories, institutions, and computing centers in the world where U.S. nuclear physicists are involved. The ESnet Review Committee's recommendation (March 1989) of a peer-to-peer connection to NSFnet is a particularly attractive way to achieve much of the required connectivity at low cost. A need for connections to other organizations, such as instrumentation and computer companies, DOE offices, and the Physical Review, is also recognized.

7. *Monitoring.* Development of software to monitor network performance is of importance. The Panel notes that faults in ESnet/MFEnet existed for a long time without understanding. Any user, and especially site managers, should have access to the management-monitoring tools that show utilization, saturation, location of bottlenecks, and down times of any portion of the network.

8. *Security.* There needs to be a procedure for handling security problems, with clear assignment of responsibilities. Such a procedure is currently being drafted by the ESnet DECnet Working Group (EDWG) for the DECnet protocol. The other protocols used need similar attention.

9. *Availability.* The Nuclear Physics program is highly dependent on a stable network and will become even more dependent in the future. The following requirements are essential and reasonable:

- >99% up time for connectivity between any two nodes
- mean time between failures greater than one month
- mean time to repair typically less than two hours. Down times of greater than half a day should not occur more than once per year.

Networking has become essential to the conduct of much of the business of Nuclear Physics research. Major interruptions in connectivity, therefore, cannot be tolerated. To insure that such disruptions do not occur, even due to natural causes, the network should be configured so that no single point of failure can interrupt connectivity between any two points on the backbone.

10. *Network Information Services.* General information about the network (purpose, organization, instructions, whom to contact, available services, etc.) should be made available.

Enabling Technology

This list contains functions that should be considered important improvements in ESnet and that would enable researchers to have their future needs satisfied as they arise. Implementation of items in this list should be driven by user needs and availability of funds.

1. *Additional Network Information Services.* Provide electronic bulletin board services for interest groups accessible from all DOE-supported nodes. Information files of interest to the Nuclear Physics community include news from each of the accelerator laboratories, current lists and status of experiments, American Physical Society/Division of Nuclear Physics (APS/DNP) newsletters and announcements, etc.

2. *Multimedia Mail/Image Transmission.* As communication in Nuclear Physics depends intimately on graphic material, graphic material as well as text should be easily included in mail.

3. *Load Balancing.* There should be no interference between interactive and file-transfer services. In particular, interactive service should not be degraded.

4. *Network-Wide File System.* A user should be able to access files on a remote host as though the files were on the local system.

5. *Bulk Data Transmission.* The network transmission speed should be capable of shipping bulk data without degrading network performance. Backbone transmission rates required would be greater than 10 Mbps. This is seen as possible 5–10 years from now.

6. *Real Time Image Transfers.* Also in the 5–10 year time frame, network planners should allow for further growth in bandwidth to permit graphics transmission (perhaps requiring interactive response) and possibly full video conferencing.

5.5.3 Current Networking Environment

The Nuclear Physics community is most involved in several WANs: BITnet, ESnet, HEPnet, MFEnet, NSFnet, and SPAN. The Nuclear Physics community has generally met its networking needs through utilizing the capabilities developed by neighboring High Energy Physics groups or services available through campus- and laboratory-wide computing facilities (such as BITnet). Although no meaningful number has been extracted for usage of networks by nuclear physicists, it has been estimated that the network traffic may be as high as 30% of that used by the High Energy Physics community. Another indication of traffic is that approximately 20 Mbytes of data are handled each month through the Los Alamos Meson Physics Facility (LAMPF) HEPnet (DECnet) connection.

Survey Summaries Concerning Networks and Nuclear Physics

DOE/NSF Nuclear Sciences Advisory Committee Instrumentation Subcommittee Survey. The following is summarized briefly from the report "A Survey of Electronics, Data Acquisition and Data Analysis in Nuclear Physics," written by D. F. Geesaman and P.D. Balamuth, using the data they collected from their September 1988 survey of DOE and NSF experimental sites. Several findings are important for assessing the WAN needs of Nuclear Physics:

1. Networking capabilities have expanded rapidly. Electronic mail via BITnet is almost universal. Remote login and database access capability are desired by everyone. To date most of this capability has been provided by HEPnet.

2. There is a widespread call for a pooling of information, talent, and resources in data-acquisition electronics, computer hardware, and computer software. This ranges from suggestions for bulletin board and information services to the call for a centralized data-acquisition group to coordinate hardware and software.

3. The community has an overwhelming investment in Digital Equipment Corporation hardware and software. Only two groups use the UNIX operating system; the rest use RT, RSX, and VMS on DEC computers. (Since the survey was made, the use of UNIX has no doubt grown with the increasing use of workstations.)

4. There has been a dramatic updating of the data-acquisition and analysis computers, with the average age being <3 years. These findings confirm that there is a growing use of and need for wide-area networking by Nuclear Physics experimentalists, that their primary access to networking has been courtesy of High Energy Physics, and that any network that is to be useful for them must accommodate the relatively young store of predominantly DEC computers.

In Section III.3.4 of the Report of the 1989 Nuclear Sciences Advisory Committee's Instrumentation Subcommittee, it is recommended that the funding agencies develop a funding mechanism whereby access to an appropriate computer network is provided to every practicing nuclear scientist.

NPPCN Survey. NPPCN members performed an informal survey of networking activities in 19 universities to better understand networking usage and requirements. A separate telephone survey was also conducted. These surveys are complemented by the site reports from different laboratories provided by members of the Panel. Several observations follow.

1. Electronic mail, file transfer, and remote terminal access are almost universal applications of networks. Remote database access, particularly QSPIRES at SLAC and the National Nuclear Data Center (NNDC) at BNL, represent a significant network application.

2. BITnet is the most commonly accessed network.

3. The university samples indicate a roughly equal split between DECnet and TCP/IP protocols and a high degree of NSFnet connectivity. There is a small MFEnet usage.

4. National Laboratories generally have access to a wide selection of networks, including BITnet, MFEnet, NSFnet, and HEPnet. The majority of networked computing resources in use by Nuclear Physics programs at laboratories are Digital Equipment Corporation VAXs running the VMS operating system.

5. Tail circuits at universities are primarily 9.6 Kbps.

6. The most common complaint about the present networking system is response time.

5.5.4 Forecast of Future Needs

Planning Driven by Experimental Programs. The manpower level associated with the DOE-funded Nuclear Physics program has not changed significantly in recent years, and this trend is expected to continue. The manner in which research is being carried out, however, is constantly changing. Experimental programs, once dominated by small groups working at local university facilities, are now more often carried out by much larger collaborations, where most of the collaborators' home institutions are located at great distances from the experimental facility. Since travel between home institution and the experiment is no longer a trivial matter, methodologies have evolved that minimize the need for frequent travel.

The construction of CEBAF, now well under way, and the proposed Relativistic Heavy Ion Collider (RHIC) at BNL, make it certain that the number of large collaborations will increase. More of the Nuclear Physics community will find itself operating in a fashion similar to that of High Energy Physics. The tools to allow management of these large groups and to allow them to operate productively while dispersed over large distances will be even more essential than they are today.

The size of the collaborations is already as great as 75 scientists for some heavy-ion experiments being carried out at this time; experiments at RHIC are expected to require manpower of over 200 scientists per experiment. The nature of these large collaborations is decidedly international, including participants from Europe and Japan. Management of groups this large dictates communications that are much more formal than those to which this community has been accustomed. In the past, Nuclear Physics experiments have had fewer collaborators than High Energy Physics ones, with a correspondingly lower number of network connections required.

These changes imply, for example, that coordination of the various aspects of the development of a large detector can no longer be handled by telephone; rather, electronic mail addressed to everyone concerned is more appropriate. Efficient operation requires that this mail be delivered within minutes, not hours or even days. This kind of response time allows for the give and take that has heretofore been possible only at a meeting. Limitations imposed by present electronic mail, such as the difficulty in transmitting graphics files, necessary for review of construction blueprints, need to be resolved to facilitate this coordination.

Software plays an increasingly important role in all but the smallest experiments. Simulations are used to aid in the design of detectors; acquisition software is required for the collection of data; monitoring software provides information on the performance of the detectors during the run; and finally, after the run, the results are obtained from the raw data with analysis software. Development of these programs is likely to be carried out by geographically dispersed groups. These people need on-line access to common libraries of code and documents.

During an experimental run, sample data may be required at remote institutions, since it is impractical to expect that the entire collaboration will be present at any one time during the run. These data are required in a quasi-real-time mode; i.e., they should be available within minutes of being acquired so that detector failures may be detected in a timely manner.

After the run, the pipeline established for the analysis of the data involves groups responsible for the construction of various detector components. Data must be available to each such group for the verification of the analysis algorithms. The results of early stages of analysis are often entries into calibration databases, containing, e.g., pedestals, offsets, or geometric relationships. These databases must be available to all groups simultaneously.

These needs, taken together, imply that every national experimental facility should either be on the network backbone or be connected to the network backbone by a tail circuit (network connection from an institution to the backbone) with a bandwidth of no less than 56 Kbps, preferably T1 speed.

Bandwidth is not the only measure of network performance; echo response determines whether a network connection is useful for interactive work. The U.S.-Europe link via satellite provided an echo response (approximately 0.5 second), which is unacceptable for this purpose: commonly accepted terminal response times are in the range 150–200 msec. A submarine cable connection is required in these cases.

Many groups have learned to work in a highly effective manner utilizing local area clusters of computers. For example, software can be developed using distributed libraries; or alternatively, copies of libraries at each site can be kept in step by network copy operations. In some cases, the purchase of expensive software, such as database management systems, makes it impractical to install it on more than one node. These groups would like to continue to operate in similar modes with widely dispersed computers providing distributed services. This kind of operation makes very large demands on network bandwidth. The ethernet that links a local area cluster has a bandwidth five times larger than the T1 backbone. Thus, in the foreseeable future, the T1 bandwidth will effectively not allow these modes of operations. With the advent of T3 backbones (45 Mbps), many of these functions will become feasible.

Planning Driven by Theoretical Programs and Data Analysis. The operating style of the Nuclear Physics community is diverse. The preceding discussion covers only the aspect seen to be most demanding of network performance. Some network operations are expected to affect every member of this community, however. For example, submission of abstracts to conferences is already being carried out by electronic mail. Program committees deal with electronic copies of abstracts, saving not only time but also enormous amounts of paper handling. Committees, such as this Panel, have learned to function with few meetings, using electronic mail, messages, and document exchange as the principal communication media.

Remote supercomputing is the mode of operation for nuclear theorists who need large amounts of computing power. The present network access to the Cray supercomputers leaves something to be desired, however. For example, large-output data files have to be transferred from MFE to a remote VAX for graphics or other analysis. A network with sufficient capacity to support a full screen editor and interactive graphics for the remote user of the supercomputer would be a significant enhancement to the computing power of theorists.

Many of these needs are addressed by the Windows concept, which partitions the interactive computing session into an application piece and a screen-management piece. The workstation at which the user is seated performs the screen-management role, behaving as a highly intelligent terminal. The application can then run on a remote machine that may have all the requisite files, licensed software, and databases. For many of the needs discussed here, this mode of operation provides the user with a highly desirable workstation interface. This mode, however, involves traffic that imposes large instantaneous demands on network bandwidth.

Planning Involving Multiprotocols, Network Backbone, and Tail Circuits. The upgrade of the backbone to T1 was eagerly awaited. This increase in bandwidth from 56 Kbps to T1 (1.3–1.5 Mbps) allows some groups to perform their research in a reasonable fashion.

Others are still limited by inadequate bandwidth of tail circuits. Estimates by the HEPnet Review Committee [1] suggested that the peak load per user is about 1400 bps; thus a 9600-bps circuit is suitable for no more than seven users. Of course, this estimate is valid for only the most rudimentary network functionality.

It has been mentioned earlier in this report that the Nuclear Physics community has a substantial investment in relatively new DEC hardware, running the VMS operating system. The most efficient connections among these machines require DECnet. However, the computer market appears to be moving rapidly toward hardware that offers dramatic increases in computing power but that, alas, does not run the VMS operating system. The interests of the Nuclear Physics community are clearly best served if the existing VAX machines can be integrated freely with this new hardware. When all network protocols are in compliance with Open Standards Interconnect (OSI), networking will no longer be an obstacle to this freedom of choice. A big step in this direction will occur when DECnet becomes OSI compliant. This event will happen when DECnet Phase V is released.

Access to networks with up-to-date capabilities is a necessity for the entire Nuclear Physics community, independent of funding sources. In particular, nuclear scientists whose activities are funded by NSF but who work at major DOE facilities must not be excluded. Peer-to-peer connections between ESnet and other national networks is necessary to ensure the efficient use of large DOE accelerators for this group. In addition, this feature could provide improved tail circuit connectivity for those who have access to the NSF regional networks.

The backbone capacity is only meaningful to those who can use it, i.e., those who are not limited by a bottleneck at the level of the tail circuit. It makes little sense to upgrade the backbone to T1 and eventually T3 while at the same time allowing some tail circuits to remain at 9600 baud. Interactive functions, even those as simple as reading mail remotely, are painfully slow when such a tail circuit is being used to service an entire campus or an even larger community. The minimum tail circuit capacity should be 56 Kbps.

Some of the functionalities discussed above can already profit from a backbone capacity of 45 Mbps. Distributed databases, transfer of moderately large files, and sampling of on-line data are some of these. The number of people in the Nuclear Physics community who avail themselves of these facilities will grow as they become more painless to use.

Based on the exponential growth in traffic observed on networks used in other disciplines, it is anticipated that a T3 backbone may be saturated during peak hours by Nuclear Physics traffic alone as early as 1995.

It is anticipated that other, as yet unused, functions will be supported by future networking capability. One of these is video conferencing. This application of networking technology could provide an acceptable substitute for collaboration meetings, allowing interplay of ideas in real time among several groups that are widely separated geographically. Video conferencing used in this way could save enormous amounts of time and travel expenses for the Nuclear Physics community.

In planning for the future, the network as a whole should be kept in mind. For one thing, the tail circuits' capacities must be commensurate with that of the backbone. Response time is also a factor that makes satellite links undesirable for many, e.g., interactive, uses. Finally, the network is only as useful as the (OSI top layer) applications that make use of them.

Reference:

1. HEPnet Review Committee, *High Energy Physics Computer Networking*, DOE/ER-0372, June 1988.

Appendix to Section 5.5: Nuclear Physics Computer Networking

The following is the executive summary of the report of the Nuclear Physics Panel on Computer Networking. Members of the Panel were Curtis Bemis, Oak Ridge National Laboratory; John Erskine, Department of Energy; Michael Franey, University of Minnesota; Douglas Greiner, Lawrence Berkeley Laboratory; Martha Hoehn (chair), Los Alamos National Laboratory; Mark Kaletka, Massachusetts Institute of Technology/Fermi National Laboratory; Micheal LeVine, Brookhaven National Laboratory; Russell Roberson, Duke University; and Lester Welch, Argonne National Laboratory/Department of Energy.

The Nuclear Physics Panel on Computer Networking (NPPCN) was formed during a period of rapid expansion in the capabilities of computer networks, including the evolution of MFEnet into the Energy Sciences Network (ESnet). Since networking has become critical to Nuclear Physics, it was necessary for the Nuclear Physics program office to take a serious look at what was happening in order to interact with the process in a positive way, rather than let the growth of computer networks proceed spontaneously to the possible disadvantage of the Nuclear Physics program. Hence, a broadly based panel of consultants, with representatives from the national laboratories, university facilities, and user groups, was appointed to cover the various ways that networks are used in the Nuclear Physics program.

The Panel was organized during the summer of 1988 and held its first meeting on September 13, 1988, in Washington, D.C. Additional meetings were held on November 15, 1988, in Washington; January 5, 1989, in Dallas/Ft. Worth; March 16–17, 1989, at Lawrence Livermore National Laboratory; and July 10–11, 1989, in Chicago. The purpose of the Panel was to review the present and projected status of computer networking within the Nuclear Physics program. The charge to the Panel was as follows:

- to review the present status of computer networking activity within the DOE Nuclear Physics program
- to identify present and future requirements of the DOE Nuclear Physics community for computer networking capabilities

- to assess the present and projected status of pertinent wide-area computer networks, and evaluate the way or ways that Nuclear Physics needs could be incorporated into the planned systems
- to identify mechanisms by which the DOE Nuclear Physics program office can facilitate communications with its research community about computer networking activities
- to determine the possible need for a continuing supply of information on networking activities. If deemed desirable, suggest organizational mechanisms for securing that information.

The use of wide-area computer-network services is a relatively new phenomenon for Nuclear Physics research, but one that is rapidly expanding. All major facilities funded by the Division of Nuclear Physics have outside user groups, and the number of nuclear physicists who participate in this mode of research is increasing as the experiments become larger and more complex. Furthermore, nuclear theorists are making increasing use of supercomputing facilities at the National Energy Research Supercomputer Center (NERSC) at LLNL and at Florida State University, with computer networks used as the primary means of access to these machines. Nuclear physicists, now more than ever, must include networking considerations in their planning, since large collaborations typically involve researchers at many U.S. institutions and in other countries. The construction of the new accelerator facility, CEBAF, and the proposed relativistic heavy-ion collider, RHIC, add further pressure for access to highly available and highly reliable wide-area networks (WANs).

The rapid emergence of computer networking as a "tool" for the research scientist can be attributed to the availability of low-cost integrated circuit technology upon which modern computing engines are built. Reduction in the cost of computing has allowed the research scientist to explore domains previously beyond reach. There is now a rapid proliferation of computer workstations, which means that most of the CPU power will be distributed among members of a research collaboration rather than at a central mainframe computer, and that data handling and network capacity will be the bottleneck.

Computing and computing engines are used in research to acquire data, to process and analyze data, to perform "model" experiments, to compare data with theoretical predictions, etc., and research in the field of nuclear physics is no exception. Computer networking for the research scientist has evolved explosively in this computing environment as access to geographically dispersed research facilities and more powerful computing resources became a requirement for the conduct of research. In High Energy Physics and increasingly in Nuclear Physics, the accelerator facilities and the researchers who use them are widely separated geographically. There is a similar story for supercomputing resources, also a "tool" of the researcher. Thus, the modern researcher in Nuclear Physics has a compelling need to exchange data and messages between members of a widely dispersed collaboration, to access experimental data at research facilities where the experiment was performed, and to have interactive access to supercomputing resources.

Findings of the Panel

The following summarizes the important conclusions reached by the NPPCN in the course of evaluating networks for nuclear physics.

1. The Panel finds that computer networking is an essential tool for the conduct of Nuclear Physics research, and that it needs to be planned for in the design of experiments along with accelerators, detectors, and other computing resources.

2. The Panel finds that the technical requirements for computer networking in the Nuclear Physics program are essentially identical to those identified for High Energy Physics. Differences occur mainly in the location of facilities, intensity of use, and the greater diversity of the Nuclear Physics program.

3. The Panel supports the development of ESnet, the computer network venture for Office of Energy Research (ER) programs, to meet most of the wide-area network needs of the DOE Nuclear Physics program.

4. The Panel encourages continuing active Nuclear Physics participation in the development, administration, and oversight of ESnet to insure that DOE Nuclear Physics needs are addressed.

5. The Panel finds that the migration of the ESnet and associated networks to international networking standards (OSI) is needed to enhance the productivity and efficiency of Nuclear Physics research efforts.

6. The Panel finds that the Nuclear Physics research community profits from a variety of networks (SPAN, NSFnet, BITnet, MFEnet, etc.) and that improved connectivity between these networks and ESnet would greatly enhance research effectiveness.

7. The Panel finds that for DOE Nuclear Physics networking needs, the existing committee structure within ER (the new Network Access Committee, ESnet Steering Committee, ESnet Site Coordinating Committee, ESnet DECnet Working Group, and HEPnet Technical Coordinating Committee) is sufficient, provided that Nuclear Physics programmatic interests are represented.

8. The Panel finds that the practice of funding tail circuits out of programmatic money, based on the needs of the individual research programs, is appropriate for the DOE Nuclear Physics program.

9. The Panel finds that many members of the DOE Nuclear Physics community, in the conduct of their research, lack a direct way of obtaining networking information and assistance in the solution of problems that arise in the use of networks.

10. The Panel finds that there is at present no adequate procedure for keeping the Nuclear Physics program office informed of the specific needs and problems of the nuclear physics research community with regard to computer networking.

In view of these findings, the Panel makes the following recommendations.

1. The Panel recommends the appointment of a Nuclear Physics Networking Coordinator, whose principal job functions are these:

   a) to act as a technical consultant for the Nuclear Physics community regarding networking questions,

   b) to communicate network news and information to the Nuclear Physics community,

   c) to monitor and calculate, wherever possible, network utilization and performance characteristics for the purpose of planning future network requirements and identifying (and resolving) current network bottlenecks, and

d) to act as advocate for the networking needs of the Nuclear Physics community.

2. The Panel recommends that an annual review and report of the status computer networking be made by a panel chaired by the Nuclear Physics Networking Coordinator.

At present (2/6/91), the Network Coordinator for Nuclear Physics has not been appointed. The HEPnet Network Manager may be able to meet the requirements for this position as well as to provide service to the High Energy Physics community.

5.6 Health and Environmental Research

The program of the Office of Health and Environmental Research (OHER) has two main missions: (1) to develop the knowledge base necessary to identify, understand, and anticipate the long-term health and environmental consequences of energy use and development; and (2) to utilize the Department's unique scientific and technological capabilities to solve major scientific problems in medicine and biology.

The OHER program includes research in atmospheric, marine, and terrestrial processes; in molecular and subcellular mechanisms underlying human somatic and genetic processes and their responses to energy-related environmental toxicants; in nuclear medicine and epidemiology; and in structural biology and other areas that require the unique capabilities required to achieve programmatic success.

Success in meeting OHER's wide range of programmatic objectives requires contributions from each of the three components of the Nation's research community: the National Laboratories, the universities, and the private sector. OHER supports research in each of these sectors and encourages cooperation among them.

Like any large-scale research program, activities being conducted vary from year to year as existing projects are completed and new projects or lines of investigation open up. The following is a synopsis of the principal research activities currently under way.

5.6.1 Health Effects Research

The Health Effects Research program conducts research in two broad areas: Biological Research and General Life Sciences. The biological research subprogram is designed to yield experimental data needed to predict long-term health impacts that may be associated with energy production and use. General life sciences research provides the fundamental understanding of the properties of living systems that is necessary for interpreting and extrapolating experimental health effects data.

The biological research subprogram encompasses both radiation biology and chemical toxicology. Of particular current interest in radiation biology are the evaluation of the potential health impact of indoor radon and the development of an understanding of the effects of high-linear-energy transfer radiation, including neutrons. Other research seeks to develop and use techniques for measuring directly in humans changes in body molecules, cells, and organs that lead to adverse effects such as cancer and genetic changes. Within the area of chemical toxicology, research is ongoing to understand the multi-level general principles of chemical effects, with particular focus on cellular and molecular mechanisms. Structural biology research takes advantage of unique resources at the Department's laboratories (such as synchrotrons, reactors, and special microscopes) to provide biological scientists from many institutions the opportunity to study the structure of biological macromolecules.

Basic research in cellular and molecular biology is needed to understand how physical and chemical agents interfere with life processes. A fundamental knowledge is sought of how living systems function and how they regulate themselves, especially their protective mechanisms for preventing or mitigating the adverse effects of toxic agents.

The aim of the Human Genome Initiative is to develop the resources and technologies that are needed to describe the human genome at the molecular level. Current goals of the program are (1) development of technologies to carry out large-scale mapping and sequencing of DNA, (2) creation of physical maps made up of ordered DNA fragments for each human chromosome, and (3) development of necessary database management and analysis tools.

5.6.2 *Human Health and Assessments*

The Human Health and Assessments Division supports research directed to increasing the understanding of human health effects that may result from exposures to radiation and to energy-related chemicals, and to the beneficial applications of energy technologies.

Current research methods encompass the study, diagnosis, and treatment of diseases such as cardiopulminary disease, mental disorders, tumor localization, cancer, and metabolic disorders.

5.6.3 *Physical and Technological Research*

This subprogram is the physical sciences component of the Health and Environmental Research program. It addresses a broad range of fundamental and applied activities and often interacts with the biological and ecological components. It also involves development of experimental techniques or advanced measurement instrumentation.

The major areas addressed are physical and chemical characterization studies related to radiation and chemical dosimetry research; radiological and chemical physics studies to determine radiation interaction mechanisms and define energy transfer and degradation processes; and development of advanced instrumentation and technology for structural biology and DNA sequencing.

5.6.4 Ecological Research

The goal of this program is to provide the basic foundation for understanding and predicting behavior of energy- and weapons-related contaminants as they are mobilized and moved from their sources to terrestrial and oceanic systems with pathways back to humans, and to determine resiliency of ecosystems to these combined impacts.

To obtain this information base, comprehensive research in terrestrial and oceanic systems is conducted in different climatic regions of the United States and in offshore regions along the eastern and western United States, along with related laboratory research. For maximum cost and scientific effectiveness, most of the research is organized into multidisciplinary, multi-institutional programs focused on regional or defined ecosystems areas. The programs cover all environmental aspects of the physical and biological sciences.

5.6.5 Atmospheric and Climate Research

Research in the atmospheric sciences concentrates on the transport and transformation of energy-related emissions in the atmosphere. Climate research is conducted as part of the Carbon Dioxide Research Program and seeks to develop a scientific understanding of the relationship between emissions of carbon dioxide (CO_2) and other greenhouse gases, climate change, and effects on ecosystems.

Emphasis of the current research is on the General Circulation Models (GCMs), which are the principal predictive tools of climate change. Other research addresses the fluxes of CO_2 among major sources and sinks of the carbon cycle. Also, the photosynthesis, physiology, and water use of plants and trees are studied to provide a basis for predicting the responses of vegetation to rising atmospheric CO_2.

5.6.6 Facilities

The National Laboratories and other onsite specialized laboratories supported by the Biological and Environmental Research budget represent an impressive national resource of skilled investigators, advances in major scientific instrumentation, and exceptional, specialized, often unique, research facilities. They provide the ability to address the diversity of research issues previously described while resulting in additional benefits to the scientific community at large, including, often, their availability as "user facilities" by qualified investigators from the academic and private sector.

Facilities playing important roles for addressing Health and Environmental Research (HER) include, but are not limited to, Los Alamos National Laboratory, Brookhaven National Laboratory, and Lawrence Livermore National Laboratory. Table 5-4 lists the major OHER DOE sites.

Table 5-4. Major OHER DOE Sites.

Brookhaven National Laboratory
Environmental Measurement Laboratory (New York)
Laboratory for Radiological Biology and Environmental Health (at Oak Ridge National Laboratory)
Lawrence Berkeley Laboratory
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Pacific Northwest Laboratory
Oak Ridge National Laboratory
Savannah River Ecological Laboratory
Stanford Synchroton Radiation Laboratory

The supercomputers and facilities at the National Energy Research Supercomputer Center (NERSC) and the Florida State University Computer Center, in addition to databases (e.g., GenBank at Los Alamos National Laboratory and the Protein Data Bank at Brookhaven National Laboratory), are utilized by HER scientists in their research.

5.6.7 Functional Requirements

HER researchers depend on computer networks primarily for the exchange of electronic mail and empirical data among colleagues. In addition, many of the researchers located at geographically dispersed DOE sites and at universities require interactive access to the supercomputer facilities at NERSC. Scientists from three OHER Divisions have received Grand Challenge awards, large blocks of supercomputer time at NERSC, to attack important fundamental problems. To take advantage of this opportunity and to use it effectively, OHER scientists need to have very-high-speed access to the supercomputers on which they run their modeling codes. In addition, the OHER community requires access to distributed databases of biological information, especially for genome and protein sequence and structural information. Limiting the ability of genome investigators to exchange information and to access data banks would seriously restrict progress in the Human Genome Project. Without such access it might take investigators years to find out that a sequence they were working on had been previously discovered. It would also not be possible to make comparisons among nucleic acid or protein sequences to study potential biological or pharmacological relationships, and many opportunities to produce new biotechnology products would be lost.

HER requirements span many of the network applications used by other ER programs. Computationally intensive modeling of data into three-dimensional visual models has become an important way of analyzing and understanding biological and environmental phenomena. As the models become more complex, extremely large databases are required to provide the needed data points, and sophisticated graphics are needed to develop visual models. This method of presenting data is being used in the Global Climate Change Program at Lawrence Livermore National Laboratory, and the interpretation of data presented in graphical form has been more fruitful than just having the data analyzed mathematically. Viewing graphics remotely across a wide-area network demands large amounts of network bandwidth; Global Climate modeling researchers, as well as those in the Structural Biology and Human Genome programs, would in particular benefit greatly by having available to them the bandwidth needed for remote graphical displays.

5.6.8 Current OHER Networking

Scientists supported by OHER are allocated computing time at NERSC and at the Florida State University Computing Center for their research programs. Table 5-5 lists the FY 1990 OHER supercomputing time allocations. In addition, many HER scientists make extensive use of remote databases, such as GenBank at Los Alamos National Laboratory and the Protein Data Bank at Brookhaven National Laboratory. The typical network applications, such as electronic mail, file transfer, remote job submission, and remote login, are routinely used by the HER researchers to access computational resources; in addition, the networks provide an easy and efficient method of exchanging information among colleagues. HER also takes advantage of Library access, via networks, to the Technical Information Center at Oak Ridge, TN. The HER scientists can be categorized as either being "onsite," which means they are located at a DOE facility such as the ones listed in Table 5-4, or "offsite," such as a scientist at a university who is collaborating on a DOE project and therefore accesses DOE facilities and computing centers.

Table 5-5. OHER Computer Time Distribution for FY 1991.

Division	CRUs
Health Effects Research Division	5,878
Atmospheric and Climate Research Division	9,624
Physical and Technological Research Division	3,810
Ecological Research Division	140
Human Health and Assessment Division	580
Total	20,000

The networks used by both the onsite and offsite scientists are varied and numerous. They include ESnet, BITnet, Internet, MFEnet, OMnet, and HEPnet, with an underlying use of local area networks (LANs) for access to these wide-area networks (WANs) and for providing network access at the local level to machines and personnel. Dial-up access is also utilized when deemed appropriate. It is important to realize that ESnet is being utilized by HER researchers even during their use of BITnet, HEPnet, and NSFnet (both backbone and regionals), since MFEnet and HEPnet are logical networks that use ESnet as a backbone. In addition, the "offsite" researchers that are currently located at NSFnet (Internet) and regional sites will transit (use) ESnet to reach the computational resources located at the DOE national laboratories and sites. All of the major OHER DOE laboratories and sites are currently members of the ESnet backbone.

The computer systems currently employed by the HER researchers include VAXs, PCs, Macs, IBM mainframes, Cray supercomputers, and UNIX workstations. There has been a dramatic increase in the use of UNIX workstations, which use the TCP/IP-based protocols and applications, by the HER scientific community to provide enhanced interactive control over the modeling codes they run on the supercomputers. The majority of the visualization workstations used by HER researchers for analyzing the results of their calculations are UNIX-based. With this current and continuing increase for UNIX-based workstations also comes the increased demand for distributed applications such as the Network File System (NFS), distributed windowing and graphics (X-11 Windows), and the underlying Remote Procedure Call (RPC) mechanisms that are necessary for a distributed computing environment. These distributed capabilities constitute network requirements orders of magnitude over the existing ESnet capabilities and capacity, which will need to be addressed within the planning period.

5.6.9 Future Requirements

As the Human Genome, Structural Biology, Global Climate Change, and other OHER programs evolve, the need for more widespread and sophisticated network applications by HER scientists will grow. The amount of data analyzed and sequences put into data banks has been doubling each year. As the Human Genome Initiative progresses, the quantity of data expected to be in databases in 15 years will be 100 to 1000 times that existing now. The sequencing effort will require greatly increased needs for computational power for improved data management, and analytical capabilities for communication and data transfer links between experimental, analytical, and theoretical scientists in many scientific disciplines. The HER community is widely dispersed and includes foreign collaborators. All can benefit greatly by using networks to maintain an ongoing dialogue with their colleagues, in addition to having access to health and biological and other related databases. Available networks made it possible for HER scientists to utilize, at their home base, the supercomputer time at NERSC they have won. Access to the NERSC Center, via high-speed networks, will allow the scientists to increase their productivity via interactive login capabilities, access and movement of large volumes of data, and provide the ability to remotely display sophisticated graphical representations of data.

The Global Climate Change and the Computer Hardware Advanced Mathematical and Model Physics (CHAMMP) scientists at different HER sites are already moving extremely large amounts of data between the supercomputers (on which they perform their modeling) and the workstations (with which they view and analyze the results). They have expressed a need for a minimum of T3 speeds between the supercomputers and their workstations, in addition to the supporting distributed applications software such as NFS and X-11. They project that their research could be greatly accelerated with easy-to-use gigabit network connectivity.

Although the current use of networks by many OHER programs is fairly basic, such as electronic mail, file transfer, and remote login, it is expected that the sophisticated use of the network will rise as the users become more cognizant of the network application capabilities available to them by using networks. An example of such use would be video conferencing. The Human Genome and Global Climate Change communities are widely dispersed and yet have a need for face-to-face communications at both a committee level and at a colleague level. Video conferencing is an application that could be used to augment the current communications structure by allowing more frequent personal interaction and collaboration. This is an area that OHER has expressed an interest in, and has given the indication that they would be willing to experiment with it if given the opportunity.

5.7 Scientific Computing

5.7.1 General Program Introduction

The Applied Mathematical Sciences (AMS) subprogram, in the Office of Energy Research, is chartered to support research and development of future computational resources for Department of Energy programs. The AMS subprogram was initiated at the suggestion of John von Neumann in the early 1950s to enhance understanding of the utility of the new digital computers in solving problems of nuclear weapons design. Over the past 38 years the nature of the research supported by AMS has evolved with the mission of the Office of Energy Research. The AMS program is managed by the Scientific Computing Staff, a staff office reporting to the Director of Energy Research. A thorough explanation of the AMS program can be found in *Summaries of the FY 1989 Applied Mathematical Sciences Research Program,* [1] available from the Scientific Computing Staff. Much of the information included here under the heading of General Program Introduction has been extracted from that document.

The mission of the AMS subprogram is twofold: (1) to ensure the broad range of research in the mathematical and computer sciences necessary to underpin and advance the physical and biological sciences; and (2) to manage Energy Research supercomputer facilities and the Energy Sciences network.

The components of the AMS program can be categorized into four major efforts:

1. The Analytical and Numerical Methods section encompasses a wide spectrum of activities in support of the mathematical modeling of physical systems, and includes work in applied analysis and computational mathematics.

2. The Information Analysis Techniques section encompasses a focus in the areas of geometry, statistics, and data management. The premise is that, as mathematical models become more complex and the computer systems for solving computational problems become more sophisticated and capable, facilities for human understanding of the results of large scale complex computations must lead the way.

3. The Advanced Computing Concepts section ensures that an integrated computational environment exists among the scientific modeler, the mathematician, the numerical analyst, the language designer, the compiler writer, the operating system designer, the

computer architect, and the theorist. This activity encompasses the advancement of both software methodology and high performance systems.

4. The Energy Sciences Advanced Computing section provides for the existence and operation of a supercomputer access program for all ER researchers, including the mandate to manage and operate the Energy Sciences Network (ESnet).

5.7.2 Functional Requirements

Principal Investigators of the Applied Mathematical Sciences program area depend upon computer networks for their effectiveness in scientific research. AMS requirements for a computer network infrastructure arise from the need to collaborate among geographically dispersed researchers and to access unique and remote computational data resource facilities. The basic need for collaboration and remote access demanded by AMS scientists has pushed the envelope of existing ESnet connectivity and services.

These evolving functional requirements for advancement of needed networking facilities can best be observed through the following contexts:

- Electronic mail, file transfer, and login to remote facilities is the most prevalent pattern of current ESnet usage.
- Remote mounted file systems are seen as part of the local user environment, increasing the demand for an embedded and transparent network.
- Energy Research scientists need access to unique AMS computational facilities for the purpose of understanding and advancing the effectiveness of high performance computer systems.
- Advances in windowing system technology, and the reduced cost and resulting availability of such systems, will provide ER scientists better access to remote facilities, which will place increased demand on ER networking facilities.
- Increased availability of computer imaging and graphics tools will push network demand by at least two orders of magnitude.

- Real-time interaction with remote unique facilities such as remote color video graphics is expected by ER scientists.
- Video conferencing is a recognized requirement in collaborative scientific research, and networking methodology (such as packet switched video) and availability will need to advance to support such an environment.
- Distributed computing among multiple sites is a technology that will migrate from the AMS environment into the general ER environment as AMS researchers advance the linkability of distributed processors (i.e., hypercubes) and networking bandwidth is expanded in order to support such an architecture.
- There is an urgent requirement to operate experimental equipment remotely; AMS scientists are presently investigating this requirement.

5.7.3 Current AMS Networking

AMS researchers utilize the networking facilities of ESnet, MFEnet, HEPnet, BITnet, NSFnet, and the scientific regional networks. Extensive collaboration exists among the National Laboratories and universities throughout the entire United States, extending to both Europe and Japan. Several AMS researchers also utilize and conduct research on the DARPA-supported gigabit testbed facilities.

Bandwidths vary from site to site, with National Laboratory researchers connected to the combined ESnet, MFEnet, and HEPnet architecture at 1.54 megabits per second (Mbps), with equivalent connections to a regional scientific network, and 9.6-kilobit per second (Kbps) connectivity to BITnet. Most ER-supported university researchers are connected at 1.54 Mbps to a regional scientific network, with lower-speed connections to ESnet and BITnet. Some of the AMS principal investigators at the National Laboratories are collaborating in research for advanced network protocols and applications using gigabit network bandwidths. (See Appendix E for an example.)

As the administrative and programmatic links to DOE tend to be stronger for ER researchers than to the providers of the non-DOE networks, AMS expectations are that ESnet should evolve in a manner quite supportive of ER programmatic needs.

5.7.4 Forecast of Future Needs

The evolving functional network requirements for the ER community will remain at the forefront of computational science-based research. The reader is provided here with some examples reflective of the changing needs for networking on the part of the AMS community. The reader is advised that the following list is very incomplete and is provided for example only.

As larger computational resources are provided to DOE from the ER sites, higher bandwidths will be required to transfer large programs and large amounts of data from the local site to the hosting site.

Graphics tools will make the user environment much more productive, with resulting increased demand for network bandwidth.

Collaborations are being formed among researchers funded by various government agencies, as in the case of global climate modeling. The computational infrastructure required for such problems is not defined, but the collaboration will be large and widely distributed. There exist unique computational and data resources that will place a large demand on the existing networking infrastructure.

As AMS collaborators make available both advanced computer architectures and tools to the ER scientific community, the breadth of the collaborating communities will increase, requiring a larger capacity and better connected network.

If the Federal High Performance Computing and Communications Program described by the Office of Science and Technology Policy (OSTP) is to be successful in advancing high performance computer systems, software, and facilities, it will rely as part of its foundation on increasing ubiquity and bandwidth among the expected laboratory, university, and industrial collaborators.

The user friendliness and raw computing power of the DOE supercomputing environments provide ER researchers with the ability to deal with significantly larger problems. As a result, larger amounts of data can be transferred between the supercomputer and the local site, where the data can be analyzed.

The future computational environment of ER sites is viewed in terms of a seamless environment, with a projection of networking demand increased by three orders of magnitude as the environment evolves.

Reference:

1. *Summaries of the FY 1989 Applied Mathematical Sciences Research Program*, DOE/ER-0422.

5.8 The Superconducting Super Collider Program

5.8.1 Introduction

When complete, the Superconducting Super Collider (SSC) will be the world's largest and most powerful particle accelerator. The SSC is currently in the design/construction phase and is expected to become operational in approximately eight years. It will be used primarily by the High Energy Physics (HEP) research community and will become the dominant HEP research facility.

Superconducting Super Collider Laboratory (SSCL) personnel now occupy temporary facilities in south Dallas County, Texas. The permanent campus will be located in Ellis County, Texas (south of the present leased temporary facilities). Occupation of Ellis County facilities will begin in FY 1991 and will continue over the next several years. Immediate SSC network requirements include provision for linking these two sites during this transitional period. Figure 5-4 shows the location of the temporary and permanent facilities.

The SSC computing strategy differs fundamentally from that of many of the established DOE laboratories. Computing resources will be highly distributed in geography, function, and management, rather than centralized through the use of mainframes. The power of UNIX/reduced instruction set computing (RISC) processors will be exploited both on the desktop and in "ranches" (assemblages of processors combined to address specific applications and computing needs). The level of commitment to this strategy is exemplified by SSC plans to purchase about 4000 million instructions per second (MIPS) of UNIX computing resources. Other HEP laboratories (and the research community in general) are expected to begin concentrating computing resources into UNIX platforms in the future to take advantage of lower-cost processing power.

Workstations are the norm on the desktop, providing local processing power, personal management tools, and applications. The processor farms provide resources to perform such functions as simulations for the design of the large detectors required by the SSC. This strategy places a very high emphasis on networking. In addition to connectivity, local networking provides the means for sharing essential system elements such as storage, input/output devices, and queues.

Location of SSC Facilities

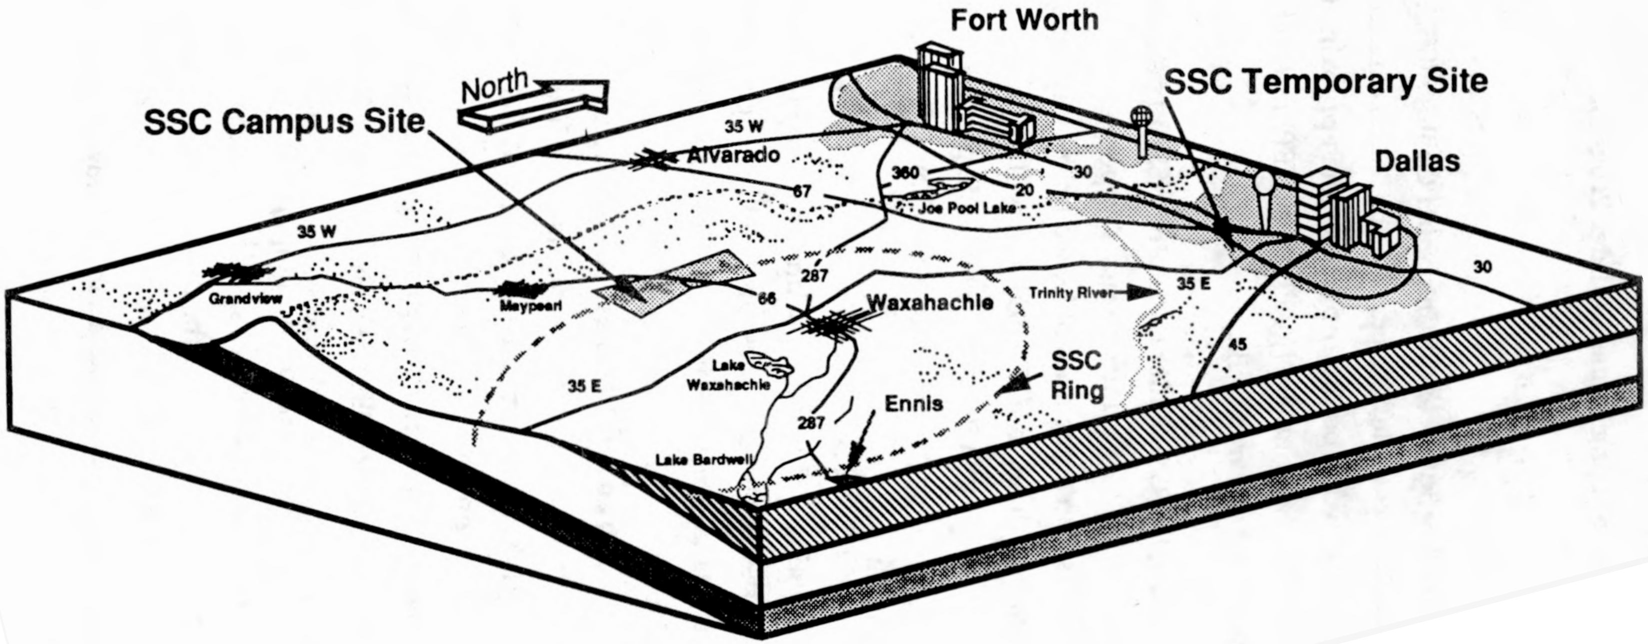

Figure 5-4. Location of temporary

At present, SSC network requirements are dominated by design- and construction-related activities. The major difference between the SSC and other HEP programs is the inclusion of connectivity with entities related to the construction of the accelerator, experimental facilities, and detectors. Good connectivity is required to existing HEP sites now, to support their participation in SSC design and development efforts. Ultimately the network requirements of the SSC will merge with those of the HEP research community, but presently the nature and content of this network activity is significantly different from future normal, research-related network traffic.

The utilization of open systems and international standards is essential for successful integration of multi-vendor UNIX systems. Consequently, it will be necessary to support the Open Systems Foundation efforts, OSI networking/government OSI protocol (GOSIP), POSIX, and any future open kernel and standards that might be developed.

5.8.2 Programmatic Functional Requirements

The SSC program needs support for a multitude of functions that require networking services. Some of these requirements have been identified in documents related to scoping the SSC project. [1] Several other documents and plans have identified networking as an essential utility for providing computing and engineering resources.

The external user network profile proposed by HEP [2] can be considered characteristic of SSC Laboratory remote user needs. The major distinction between the SSC and HEP programs stems from the fact that UNIX-driven computing systems require a greater emphasis on TCP/IP networking, while HEP relies primarily on the DECnet protocol.

Dual-protocol support (TCP/IP and DECnet) is required by the computer systems and networks in use at the SSC, but because the detector-simulation computing resources and detector collaborations are not yet established, there can be no meaningful estimates of network traffic patterns or volume.

Network Applications

Anticipated SSC Laboratory requirements include the following:

X-Windows. The computing resources identified for simulations will be used by members of detector collaborations at existing HEP sites. Efficient use of these resources will require support of X-windows and appropriate network throughput from the SSC Laboratory to the collaborator's institution.

Video Conferencing Requirements. New technology has made video conferencing a practical means of communication, and the SSC Laboratory has a keen interest in employing video conferencing as a tool in project management, for liaison with the HEP community, and in support of the detector collaborations. [3]

A proposal submitted to DOE, and subsequently approved, defines a video network designed to serve the specific needs of the SSC program. Implementation began with a pilot project that linked LBL, FNAL, and SSCL in the spring of 1990.

The real-time, circuit-switched requirements of video demand different networking resources that could impact the ESnet backbone configuration and possibly create new types of service. For this reason, some research and development investment in packet video technology and standards establishment should be considered in the best interests of all ER programs.

The SSC program wishes to address video connectivity needs within the domain of ESnet. There is no intention or desire to establish or maintain a separate SSC network to support video conferencing. Other ER programs (particularly HEP) are expected to identify additional video requirements of their own that will integrate with those of the SSC program.

Traditional Network Application Requirements. Support for the traditional network applications, such as electronic mail, file transfers, task-to-task communications, and virtual terminal support, is required for the SSC program.

It is anticipated that network queueing, printing, and naming applications will have to be provided outside the scope of ESnet. ESnet can anticipate the SSC program's use of these applications for both internal and external users.

Tail-Circuit Performance Upgrades

Support of X-windows and network-distributed file systems will necessitate more than additional network backbone bandwidth. The typical 9.6-Kbit line from the university to the ESnet backbone cannot support these more demanding applications. Upgrading could be achieved by increasing tail-circuit speeds to 56 Kbits; this would be very costly if done universally across the network. Using fractional T1 circuits could reduce the cost of upgrading and should be considered as an option.

The ESnet plan to connect directly to NSF regional networks is potentially an excellent means of providing tail-circuit performance enhancements. ESnet should carefully select regional connection points to maximize the benefits to the ER programs. Consideration should be given to increasing bandwidths within a regional network or bypassing local bottlenecks if program requirements dictate.

Network Connectivity to Industrial Affiliates

During construction of the SSC accelerator, facilities, and experimental components, connectivity to industrial entities will be required for the coordination and management of their efforts. Links to industrial entities will support many applications, including CAD/CAE, video teleconferencing, and electronic mail. These links will not be required once the work is complete.

The SSC does not wish to establish a separate wide-area network to meet these requirements. It is preferable that they be supported under the umbrella of the ESnet charter, with appropriate precautions taken to ensure adequate security for both ESnet and industrial sites.

International Networking Requirements

The SSC program will attract participants from many countries outside North America, some of which are already connected to ESnet in support of their involvement in other ER programs. However, several countries not currently connected to ESnet will need connectivity to the SSC. Several international collaborators will wish to use video conferencing in addition to the traditional network applications.

Because of the additional activity and use of more demanding applications associated with the SSC, line upgrades may be needed to specific countries to meet SSC requirements. Ideally, SSC international requirements will be factored into ESnet, along with those of the other ER programs.

Current SSC Wide-Area Networking

Virtually all SSC networking connectivity is provided by ESnet. Esnet maintains four T1 circuits to the SSC temporary site, linking it to FNAL, FSU, ORNL, and UT-Austin. These circuits currently provide adequate throughput and redundancy. Two hundred fifty-six Kbytes of the T1 circuit to FNAL will be used to support video conferencing by linking the SSC and FNAL video branch exchanges.

The SSC Laboratory independently maintains a T1 circuit to UT-Dallas. This circuit provides connectivity to the two Texas regional NSF networks, Sesquinet and the Texas Higher Educational Network. Both DECnet and TCP/IP are supported on these networks. BITnet connectivity for the Laboratory is maintained over a DECnet link on this circuit to a UT-Dallas VAX system.

The SSC does not currently maintain any tail circuits. All tail circuits associated with SSC activities are maintained by other HEP laboratories in support of their participation in the HEP program. [4]

Summary of SSC Program Requirements

The following applications/needs use ESnet resources for support of the SSC program:

- TCP/IP and DECnet
- X-windows
- network-wide file systems
- tail-circuit performance upgrading
- video conferencing
- file transfer

- virtual terminal
- electronic mail
- task-to-task communications
- connectivity to industrial affiliates
- international networking requirements.

5.8.3 Forecast of Future Needs

Many factors that determine the profile of SSCL networking needs cannot be readily predicted. The construction of the accelerator, facilities, and detectors over the next eight years is driven by funding, which is determined on a year-to-year basis. This yearly uncertainty greatly complicates SSC long-term planning.

The anticipated use of video conferencing, windowing, and network-wide file distribution will demand the next quanta of performance. The success of the SSC program will be related directly to the creation of, access to, and processing of experimental data. Even prior to the construction of the accelerator, simulated particle physics data are generated, passed through modeled detectors and accelerator components, and analyzed using computers. These steps are necessary for the successful design of highly complex detectors and accelerators.

High-performance networks must be available for use through the design and construction of the SSC and into the eventual operation of the scientific program. It can confidently be stated that gigabit networking is absolutely necessary to support the SSC program in the near future.

Needs for Period FY 91–92

Video Conferencing. The initial connectivity needed to support video conferencing was identified in a proposal submitted to the DOE. [5] Included in the proposal were requirements to establish video at 384 Kbytes from the SSC facility to LBL, FNAL, BNL, and DOE/OSSC in Germantown for administrative support and magnet research and development. Additional connectivity in support of detector collaborations at up to 30 locations (15 each in FY 91 and FY 92) was identified.

The video conferencing pilot project was extremely successful; it will be continued and expanded. At this time (November 1990), 11 specific locations have been identified for video conferencing support, including KEK in Japan and INFN in Italy.

Due to the immaturity of the networking aspects of video conferencing systems, it is expected that some temporary work-arounds may be necessary to accommodate video applications within the ESnet framework.

Detector Simulation Facility. In FY 91, procurement of 1000 MIPS of computing resources will proceed in support of detector simulation. By the end of FY 92, there will be a total of 4000 MIPS purchased and functioning. At this time it is difficult to estimate the networking requirements this computing facility will generate. The use of X-windows and the need to move large files through the network will necessitate some selective form of tail-circuit upgrading if this facility is to be used effectively.

Networking to Industrial Affiliations. During the current design/construction phase, the SSC program is expected to require serveral connections to industrial affiliates in support of both networking and video conferencing applications.

Network Backbone. The total network traffic associated with the support of video and the detector simulation facility may mandate upgrades in the ESnet backbone. Assuming other programs will develop similar requirements, selective circuit installations in the T3 bandwidth range may be required.

Beyond FY 92–98

In this period the SSC accelerator, facilities, and detectors will be built, and the experimental program will begin full operation. The need for considerable network support is anticipated through this period.

As major accelerator elements become operational, test beams will be used to verify the operational status of detector and beam monitoring components. Some parallels between the HEP and SSC program requirements will be drawn as similar aspects are engaged.

Beyond FY 92, full utilization of the permanent campus facilities and the movement of all personnel and equipment to Ellis County will occur. At an appropriate time, all ESnet connections will have to be reterminated at the Ellis County campus location.

Utilization of network-wide distributed databases and file systems may become a vital element in support of future physics data analysis. Each detector collaboration will consist of several dozen institutions, geographically dispersed, and including international participants. Sharing and coordination of data among this multitude of entities is a major concern that can be addressed, in part, by a network-distributed file system.

Widely distributed databases (and database engines) may become the means by which future experimental data are analyzed, creating a new network application that will have expanded (and different) network requirements.

5.8.4 Research and Network Development

From the perspective of the SSC program, investment in research and development in the following areas is highly desirable:

- packet video
- network-distributed databases
- tail-circuit performance enhancements
- OSI migration
- wide-area gigabit networking

Esnet should be funded at a level that will allow for addressing future applications as well as providing for continuation of current networking capabilities.

5.8.5 Closing Remarks

The current ESnet backbone configuration adequately addresses SSC program networking needs. However, by the end of FY 91, operation of new major computing facilities and utilization of new network-intensive applications will begin, creating additional stress to the installed ESnet backbone circuits and saturating many existing tail circuits.

SSC experience with ESnet has been very good, both in network availability and in the responsiveness of the ESnet staff. The maintenance of adequate staffing levels as ESnet grows is extremely important.

As important as gigabit networking seems to be in the future, it is imperative that ESnet retain adequate control of resources to respond to the demands of individual programs. If programmatic control is lost for the sake of participating in a grand network venture, ESnet may not be able to serve SSC needs. DOE/OER should work to ensure appropriate isolation from any initiatives that would, in effect, supersede ESnet and reduce the ability of ESnet to respond to individual programmatic requirements.

References:

1. *Report Of The Task Force on Computing For The SSC*, M. Gilchriese, 12/88.

2. 50 Kb for 10% of the working day plus 20-Mbyte file transfer per user.

3. *Preliminary Report on SSC Video Teleconferencing Needs For Detector Collaborations*, Rev. 1, G. Chartrand, L Cormell, G. Yost, 6/25/90.

4. See HEP Program Requirements for ESnet.

5. *Proposal for Video Teleconferencing and Associated Communications*, G. Chartrand, 3/15/90.

Current ESCC Activities 6

Section 6: Current ESCC Activities

6.1 Introduction

The ESnet Site Coordinating Committee (ESCC) provides a forum for a three-way dialogue between (1) the ESnet Steering Committee (ESSC) for policy and guidance; (2) individual sites for functionality, activity, and resources; and (3) the National Energy Research Supercomputer Center (NERSC) ESnet Manager and staff for backbone and resources, including network services. The membership of ESCC is composed of representatives from each of the major ESnet sites. ESCC has been meeting three times a year during its formative period. Beginning in 1991, it will only meet twice a year, because the ESnet backbone is functioning well and the committee members have established good working relationships.

ESCC carries on activities of the committee as a whole, sets up task forces for short term projects, and has working groups for issues needing long term consideration. In addition, committee members interact daily with the NERSC ESnet staff to take care of site issues relative to keeping ESnet functioning at a high level of performance. Examples include coordinating hardware and software upgrades, providing local interfaces to the communication companies providing the links, and monitoring performance from the sites' points of view.

6.2 ESCC Task Forces and Working Groups

Modern networks do not live in a vacuum; rather, they have gateways to other networks. Thus, several of the ESCC task forces and working groups include members from the networking community outside of ESnet. Table 6-1 lists current and recently active task forces. The discussion of the current and recently active task forces is followed by sections on each of the working groups.

The IP Routing ESCC Task Force reviewed the IP Routing document prepared by the NERSC ESnet staff. The result was that, when IP routing was turned on in ESnet in the spring of 1990, everything worked quite well. All of the critical issues were known and agreed on by both the ESnet staff and ESCC. This document is now being used as a template for IP routing issues beyond ESnet.

Table 6-1. Current and recently active task forces.

Task Force	Chair	Status
IP Routing	Larry Amiot (ANL)	Task complete
DECnet Phase IV Routing	Les Cottrell (SLAC)	Task complete
Network Access	George Rabinowitz (BNL)	Task in progress
OSI Transition	Tony Hain (NERSC)	Task in progress
MFEnet Transition	Committee as a Whole	Task in progress
Resources and Education	Roy Whitney (CEBAF)	Task in progress

The ESnet DECnet Working Group (EDWG) prepared the DECnet Phase IV Routing Document. It is important to note that, the way DECnet Phase IV works, the ESnet DECnet interacts directly with routing issues with the worldwide research DECnet. It is also important to note that this worldwide DECnet was fully operational prior to moving much of the U.S. Energy Sciences DECnet onto ESnet. This document elegantly integrates all of these aspects and is used as the basis for guidance in keeping operational the very well functioning DECnet. The result was that when DECnet routing was turned on in ESnet in the late fall of 1989, everything worked very smoothly.

The Network Access Task Force has recently been formed to address issues relative to access to ESnet. The Task Force is preparing a document for approval by ESSC and ESCC to be used as guidance in the day to day work of managing this aspect of the network.

Issues relative to the addition of OSI protocols to ESnet are being addressed by the OSI Task Force. The study of these issues is in progress.

ESCC has interfaced with the ESnet staff on the migration of MFEnet I to MFEnet II. There was consideration of forming an ESnet MFEnet Working Group, but it was felt the ESCC as a whole was adequately interacting with the development of MFEnet II.

ESnet is now providing greatly improved access to many resources. Examples include CPUs, databases, bulletin boards, and technical expertise. ESCC is in the process of assembling a report/database of these resources on the network.

In the area of education, ESCC is working on a Network Users Guide. The intent is to have both a hard copy and on-line document to ease the path of learning how to access the many resources now appearing on ESnet.

6.3 ESnet DECnet Working Group (Chair: Phil DeMar, FNAL)

The ESnet DECnet Working Group (EDWG) provides planning, documentation, review, and site-level management for DECnet issues, including addressing, performance, security, and Phase V transition planning. The membership of EDWG consists of individuals well versed in the complexities of wide area networks (WANs) involving DECnet. The membership of EDWG is subject to approval by ESCC. The chair of EDWG is chosen by EDWG, subject to approval by ESCC. EDWG was a major factor in successfully bringing the pre-ESnet HEPnet and DECnet on to the ESnet backbone. During this process, EDWG created the DECnet Phase IV Routing Document, which includes addressing, interagency and Big Four (HEPnet, Space Physics Analysis Network (SPAN), Euro-HEPnet, and Euro-SPAN) agreements, circuit costs and routing, site issues, and management. EDWG has also prepared an ESnet/DECnet Security Policy, Procedures, and Guidelines document, which is currently under review by the ESCC Security Working Group.

EDWG is currently working on the details of a DECnet Phase V Transition Plan. Roughly, the plan is to bring up DECnet Phase V first on the ESnet backbone. Then, area by DECnet Phase IV area, the sites will transition to DECnet Phase V.

The working group will exist until there are no more DECnet-specific issues requiring ESCC attention.

6.4 ESnet Security Working Group (Chair: Mark Kaletka, FNAL)

The ESnet Security Working Group (ESWG) coordinates, monitors, documents, and reviews ESnet and associated site security issues, responses, and procedures. At the present time ESWG is reviewing the EDWG ESnet/DECnet Security Policy, Procedures, and Guidelines document. ESWG will then create a "generic" ESnet security document. ESWG may create a separate UNIX/IP security document. Note: these documents refer more to end node and site security issues. "Network security," while important, is a relatively minor issue. The usual analogy is that you lock access to your house, not to the freeway; you monitor the freeway. The major issue is the security of all resources on the network, including of course the network itself.

6.5 ESCC on Video Conferencing

Video conferencing is anticipated to be a major network activity in the near future. ESCC is planning to have an ESnet Video Conferencing Working Group as soon as the technology and management issues stabilize.

6.6 ESnet Backbone Sites—Total Network Usage

Table 6-2 is a report by ESCC members of current total network usage. This usage includes all of the networks, tail circuits, etc., connected to the site. For DOE sites, the ESnet fraction of the traffic ranges from fifty to ninety percent of the total, depending mostly on the number of dedicated tail circuits connected to the site. These numbers are representative for May 1990.

Table 6-2. Total Network Usage, ESnet Backbone Sites.

Site	Mbytes/day	
	In	Out
Argonne National Laboratory	78	59
Brookhaven National Laboratory	810	810
Continuous Electron Beam Accelerator Facility	30	25
California Institute of Technology	695	695
Fermi National Accelerator Laboratory	308	685
Florida State University	a	a
General Atomics (San Diego)	280	280
Los Alamos National Laboratory	180	160
Lawrence Berkeley Laboratory	188	170
Lawrence Livermore National Laboratory	100	100
Massachusetts Institute of Technology	60	75
National Energy Research Supercomputer Center	50	600
New York University	a	a
Oak Ridge National Laboratory	a	a
Pacific Northwest Laboratory	5	5
Princeton Plasma Physics Laboratory	130	110
Stanford Linear Accelerator	600	600
Superconducting Super Collider	55	52
University of California at Los Angeles	a	a
University of Texas (Austin)	a	a

[a]Not determined.

Future Needs 7

Section 7: Future Needs

7.1 Introduction

During FY 1985, the U.S. Congress mandated a study of computer networking requirements of the U.S. research community. [1] The preparation of such a far reaching study, together with the growing widespread remote access to supercomputers through Federal supercomputer programs, such as those in the DOE and NSF, focused the attention of the U.S. research community on the need for capable, high speed computer networks.

Within DOE, computer networks had already been used heavily for specific applications and programs, primarily supporting Fusion and High Energy Physics, but these networks were mostly incompatible and lacking in needed capability. Because of these limitations and significant increases in networking requirements, the Energy Research community endorsed a proposal to create the Energy Sciences Network. ESnet has been developed to be compatible with existing network requirements while providing connectivity and interoperability to other Federal research networks, in addition to providing a nondisruptive transition path to emerging international network standards.

ESnet is the vehicle through which the Energy Research (ER) community has become a full partner in the Internet community of computer networks and through which the ER community will participate in the proposed National Research and Education Network (NREN).

ESnet currently provides network access for many ER collaborators through collaboration with NSFnet (and its regional networks), as was proposed in the earlier ESnet Program Plan (June 1987). [2] This cooperative arrangement of networks is generally considered to be the precursor to NREN.

7.2 ESnet Future Requirements

ESnet was developed after careful definition and documentation of user requirements from all of the Energy Research program areas. This user-requirements orientation led to an ESnet implementation that departed from other Internet implementations. Most notably, ESnet was developed to be a multiprotocol network to preserve Energy Research user-applications investments in existing network implementations, while allowing full interoperability with the Internet. Recently, the Internet Advisory Board (IAB) has also adopted the idea of incorporating a multiprotocol approach to allow a nondisruptive transition to international standards.

7.2.1 OSI Integration into ESnet, Coexistence and Interoperability

ESnet interconnects the local area networks of the ER facilities and provides services for both TCP/IP and DECnet. These protocol families will have different migration paths toward Open Systems Interconnect (OSI) protocols but will converge at a common end point. While it is necessary to recognize the need to integrate OSI into all networks, the need to continue providing existing network services to users must be recognized as a higher level commitment. The OSI infrastructure is missing many key elements, such as interdomain routing, directory/name service, and network management. As a result, there will be coexistence and potentially difficult interoperability among OSI, TCP/IP, and DECnet protocols for the indefinite future. These issues will be resolved in coordination with the Federal Networking Council (FNC) OSI Planning Group (FOPG). Transition from DECnet to OSI will have a path provided by DEC, and specific details of a solution are being worked on by EDWG. Despite the current difficulties, there is momentum behind the OSI development, and it is anticipated to be in wide use eventually. Rapid transition from TCP/IP to OSI is not likely to occur; rather, it is anticipated that an extended transition period will take place during which multiple protocol suites will be in widespread use. Some of the most optimistic speculation has been that OSI will be in predominant use five years from now, while the more pessimistic speculation is that it may take another ten years.

The integration and transition of any new protocols into any large established network base will require adjustments or fine tuning of both the protocols and the supported applications. These effects can normally be expected for changes of this magnitude. Therefore it is necessary to identify the issues and potential problems that will arise as the new OSI infrastructure is integrated into existing networks. These issues are considered an important part of the initial planning stages for the OSI migration of ESnet.

The ER community is committed to the transition to the Government OSI Profile (GOSIP) protocols and will adopt them for use when they become commercially available and are cost effective. The ER community has many representatives on the

DOE GOSIP Transition Working Group, the National Institute of Standards and Technology (NIST) OSI Implementors Workshop (OIW), and the Internet Engineering Task Force (IETF), with the intent of advancing the use of OSI/GOSIP.

7.2.2 OSI Transition

The inclusion of OSI into the ESnet infrastructure will require several integrated stages over a period of several years. The following milestones are anticipated in the transition.

Stage one: The goal of stage one is to provide for and experiment with early versions of the OSI/GOSIP protocols (1990–?).

- GOSIP version 1.0 effective (few products available)
- DECnet Phase V router field test
- OSI applications over TCP/IP (ISODE)
- development of other aspects of OSI

Stage two: After experience is gained with early versions of GOSIP in stage one, it will be necessary to introduce the OSI protocols into the production network. This will range from the incorporation of multiprotocol routers to GOSIP↔TCP/IP application gateways (1991–?).

- GOSIP version 2.0 effective (products limited in availability)
- multiprotocol routers
- multiprotocol end systems (limited)
- encapsulation of International Standards Organization (ISO) 8473 over IP
- application layer gateways (electronic mail, file transfer, remote login)

- transport service bridges
- standard transport service interface

Stage three: This stage will provide for extensive coexistence and application interoperability between OSI/GOSIP and other protocols (1995–?).

- multiprotocol routers
- multiprotocol end systems (limited)
- transparent application layer gateways (limited to electronic mail, file transfer, and remote login)
- extended directory services
- new applications over OSI
- new routers optimized for OSI

Stage four: This stage will provide for coexistence, application interoperability, and ubiquitous OSI (2000?).

- encapsulation of IP over ISO 8473

Table 7-1 summarizes the current OSI situation.

7.2.3 *Programmatic Requirements*

The Energy Research programs have documented their requirements and presented them in Section 5. Several of the most demanding functional requirements are discussed here.

Multiprotocol Support. The Energy Research community will have a continuing need for multiprotocol support throughout the initial planning period. Even though current procurement trends indicate a continued transition to UNIX systems (and TCP/IP network protocol usage), a wholesale conversion to these systems is impractical and too

Table 7-1. Current OSI situation.

Layer/Issue	Current Status	Responsible Group/person
Physical layer	none	
Data link layer		
OSI and IP over Ethernet/802.2,3,4	done	
OSI and IP over HDLC/LAPB	in progress	
OSI and IP over X.25	in progress	
OSI and IP over PPP	work needed	
Network layer		
Routing protocols		
IS-IS for OSI only	in progress	ANSI/ISO
IS-IS for OSI and IP (dual)	experimental	IS-IS WG
Parallel stacks	in progress	
NSAP format	done	
NSAP guidelines	in progress	NSAP WG/NIST
CLNP, ES-IS, IP, etc.	done	
Interdomain routing	in progress	many efforts
ISO 8473 echo	in progress	ANSI/ISO
CLNS/CONS interoperating	in progress	ANSI/ISO
Transport layer		
Transport service bridges	work needed	IETF
Session and presentation	done	

Table 7-1 (continued)

Layer/Issue	Current Status	Responsible Group/person
Application gateways		
Electronic mail		
X.400(84) <-> RFC 987	experimental	IEFT
X.400(88) <-> RFC 1138	experimental	IEFT
Virtual terminal (VTP-TELNET)	experimental	commercial
File transfer	in progress	
Directory services		
How to set up name space	in progress	NIST
Multistack protocol query	work needed	X.500 wg/IETF
X.500/DNS interoperating	work needed	X.500 wg/IETF
Security and access control		
For IP-only environments	work needed	
For OSI-only environments	work needed	
Specific to dual environments	in progress	
Network management	in progress	
Mixed stack		
ISODE	in progress	Marshall Rose
X.400 naming and routing	in progress	NIST OIW

expensive to be accomplished within the next three years. During this time, OSI standard protocols will be incorporated into the ESnet network layer in coexistence with those currently in use (e.g., TCP/IP, X.25, and DECnet).

Video Teleconferencing. In recent years, the cost of video technology (using compression techniques for video transmission) has dropped drastically. Additionally, the future development of packetized video will provide a broad distribution means to support video within existing networks. These events have spurred great interest in using video for research collaboration support. In fact, a pilot project for video conferencing has been initiated within ESnet by the HEP and SSC programs. This video pilot project has received positive evaluations from all involved.

The benefits of video are even more evident as one examines the logistical problems of multinational collaborations in Energy Research programs. Video will enable these collaborators to meet weekly (or even daily), which would not otherwise be possible. This will greatly enhance the productivity of international collaborations and has the potential to minimize travel costs and overhead. These benefits will also be accrued by geographically distributed collaborations contained within the U.S. boundaries.

Windowing Technology. The use of windowing technologies and other software tools has enabled more facile and productive use of remote computing and control resources. These tools "extend" the capability of a user's local workstation to remote computing resources, facilitating simultaneous access to multiple sources of information, so that many sources of information can be used more productively in the investigation of complex problems.

Present windowing technology allows users to access system services distributed over a number of different processors. The future computational environment (driven by advanced windowing technology and applications) will provide a more seamless view from the user's perspective with little regard for the actual hardware and will virtually eliminate the need for procedural software. The use of these future tools in a distributed environment will most likely increase the network load by up to three orders of magnitude as the technology evolves and becomes widely used.

Network Distributed Tools and New Network Applications. There is a growing requirement for distributed design, programming, and debugging tools. Remote-procedure-call functionality has recently been implemented on the supercomputer systems at the National Energy Research Supercomputer Center. As this functionality achieves more widespread use, it is anticipated that distributed processes will place a heavy demand on ESnet resources.

For large distributed systems that will be remotely managed, it will be important to have the ability to analyze system performance with regard to parallel performance, individual system node processes, and throughput. This ability is especially needed for the evaluation of prototype 100-gigaops and teraops systems that now require nonbroadcast network topologies (making it difficult to analyze application-level network traffic). Support of these systems will place even more demands on network bandwidth (possibly another order of magnitude) in the FY 93–95 period.

Proposals for the next generation of large experiments in Fusion (for the International Thermonuclear Experimental Reactor) and in High Energy Physics (for the Superconducting Super Collider) include requirements for distributed control and/or monitoring of their experiments. Presently, such requirements demand high-performance local area networks. Extending these requirements into wide area networking (e.g., ESnet) in the FY 95–98 period will require bandwidths in the 100-megabit range.

Visualization. Recently, the use of supercomputers has proven the validity of visualization for allowing new understandings of various physical phenomena. One of the most prominent examples of the utility of visualization has been demonstrated by Dr. L. Smarr, Director, National Center for Supercomputer Applications. [3] His work, using visualization of the output (as opposed to volumes of printed data or thousands of graphs) has led to a "new" understanding of storm physics, despite the fact that the model's complexity has changed very little in its supercomputer implementation.

The Energy Research community has many similar requirements for visualization, in Global Climate Change Research and other ER applications (see Figure 7-1). Current visualization implementations require up to 800 megabits/second output capability. For example, LBL is currently developing a visualization capability for a remote medical imaging application that has been shown to require bandwidths in this range. This application is described in Appendix E. Wide-area-network implementations that

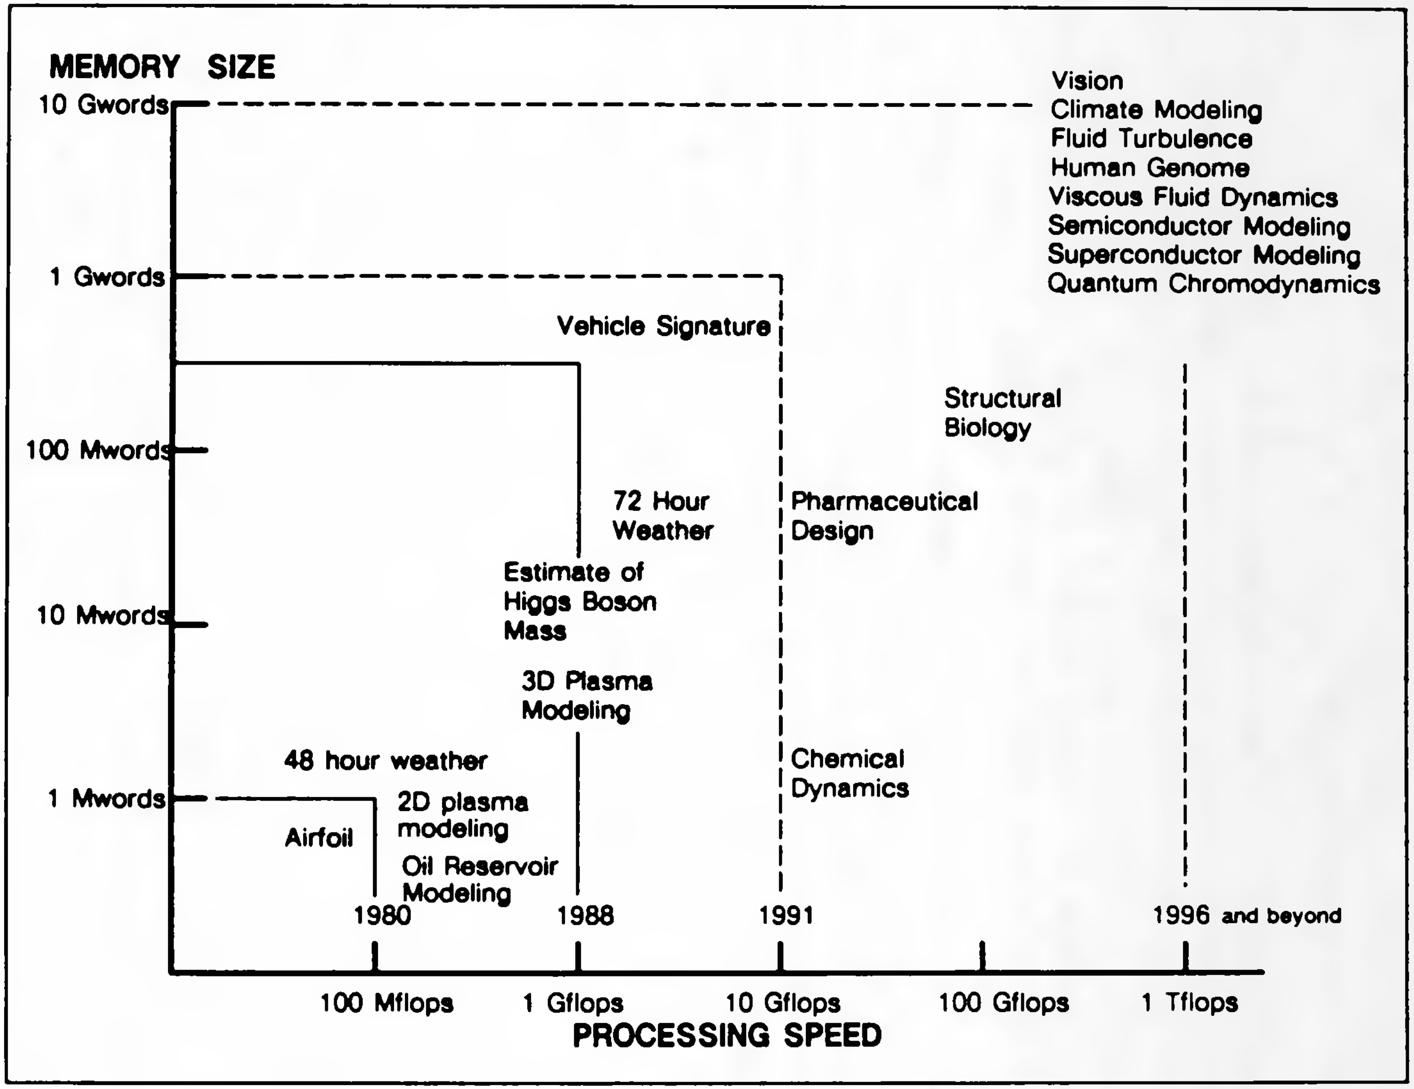

Figure 7-1. Some Grand Challenges and their projected computational requirements. (Note: one word equals 64 bits.)

require bandwidths less than 800 megabytes for visualization will be investigated for use within ESnet. However, it is anticipated that visualization techniques will generally require at least one gigabit/second, as demonstrated by the NASA Remote Interactive Particle Tracer workstation software.

Response Time. While from the user's perspective, response time is a very important parameter, bit rate is usually referenced because it is often the most important component of response time over a computer network. Ideally response time should be independent of whether the communicating end systems are a through backplane, a local area network, or a continent away through a wide-area net. This is most applicable for work group or client/server-style distributed computing.

Figure 7-2 [4] quantifies the needed response time for several types of computing operations as a function of the size of the "work units" often transferred for that type of computing. For instance, terminal-timesharing computing (as characterized by "dumb terminal" transfers) need response times between 0.1 and 1.0 second. Client-server (or work-group computing) needs are characterized by remote procedure calls (RPCs), demand paging, swapping, and interobject communication (IOC), requiring millisecond response times. Local area networks that have single-hop connectionless (CL) protocols can achieve these needed response times today. However, there is a growing need to achieve such responses over wide area networks.

ESnet Response. To meet these (and other) requirements, ESnet plans to begin the transition of its current T1 network to T3 technology (45 megabits/second) within FY 92–93 and to begin testing gigabit networking capabilities as they become available. Several DOE laboratories already participate in the DARPA gigabit network research testbeds.

7.3 ESnet Relationship to NREN

The DOE contribution to the High Performance Computing and Communications (HPCC) Program, of which NREN is a major component, is to evolve ESnet to become an integral and fully compatible component of NREN. HPCC is a balanced computing research proposal that includes research in high performance computing systems, software technology, and computer networks, as well as funding for human resources

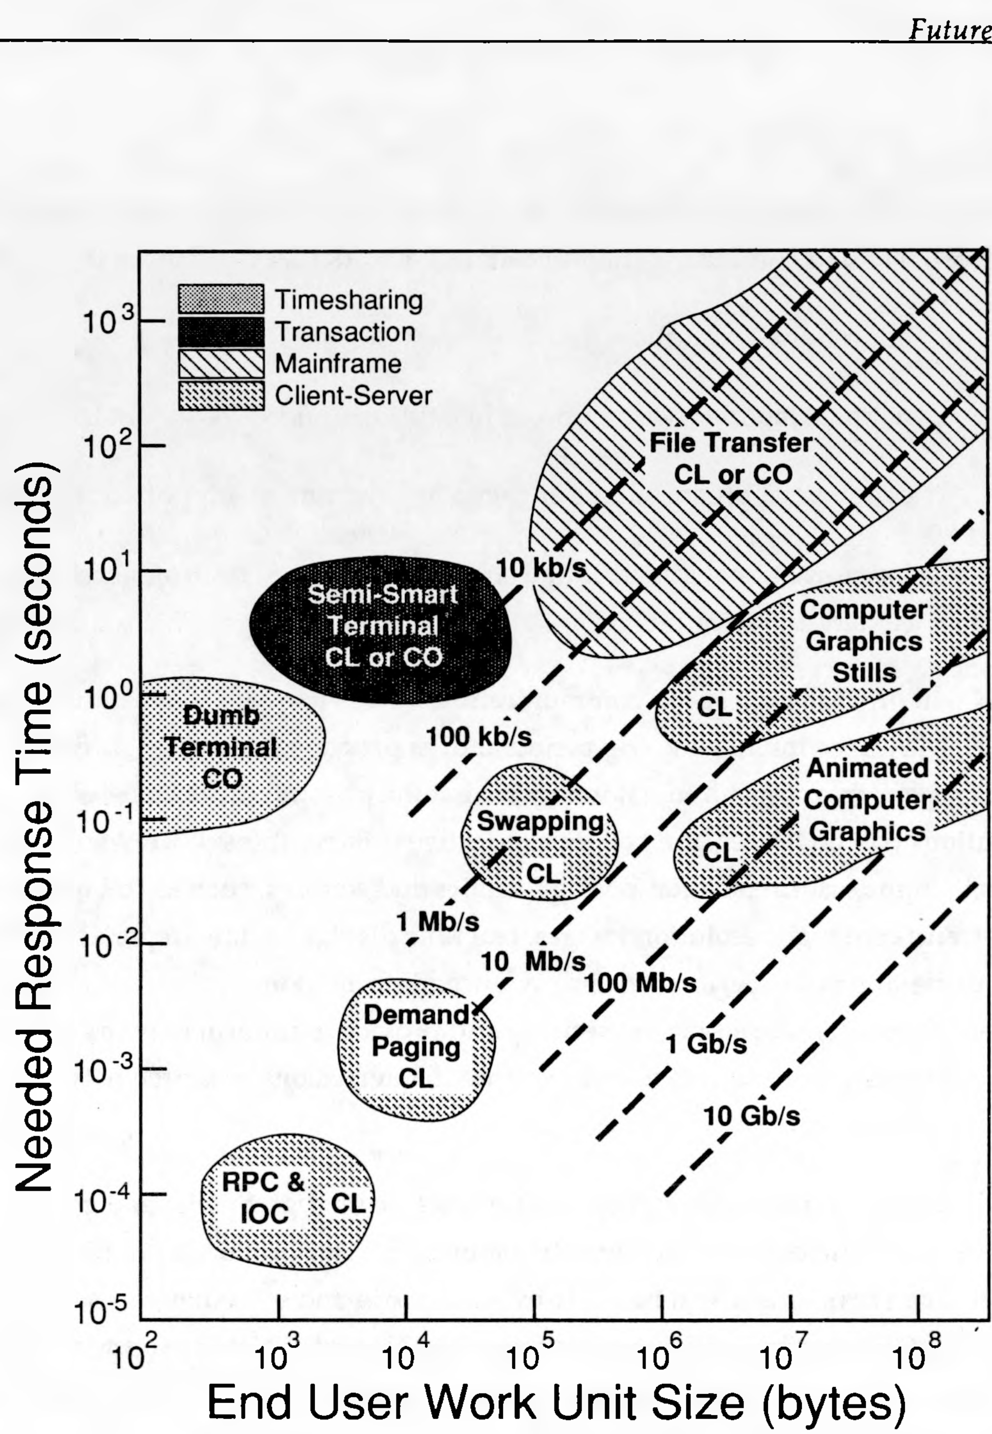

Figure 7-2. Future Computing Response Time Requirements.

and NREN. NREN, as proposed (see Figure 7-3 for the staged NREN implementation schedule), will be a computer communications network that interconnects

- educational institutions
- national laboratories, nonprofit institutions, and government facilities
- commercial organizations engaged in government-supported research
- unique national scientific and scholarly resources, such as supercomputer centers, libraries, etc.

NREN will provide high speed communications access to over 1300 institutions across the U.S. within the initial planning period and, as proposed, will offer sufficient capacity, performance, and functionality so that the physical distance between institutions will no longer be a barrier to effective collaboration. NREN will support access to high performance computing facilities and services, such as full motion video, rapid transfer of high resolution images, real time display of time-dependent graphics, remote operation of experiments, and advanced information sharing and exchange. NREN will incorporate advanced security and provide a uniform network interface to domestic users as well as a standard interface to international research networks.

NREN proposes to achieve economy of scale by serving many Federal agencies, industrial R&D centers, and university campuses. Although the Federal government will provide a substantial direct investment in NREN, it is important to recognize that a large indirect investment will be made by institutions and other networks that will connect to NREN. The DOE contribution to NREN is for the direct support and participation of the DOE laboratories and other DOE-funded research facilities attached to NREN, per the existing Federal budget process.

The ESnet Steering Committee has reviewed the NREN proposal and has recommended that DOE advocate a management structure for NREN that parallels that of ESnet to ensure a requirements-driven approach with strong user involvement and site coordination.

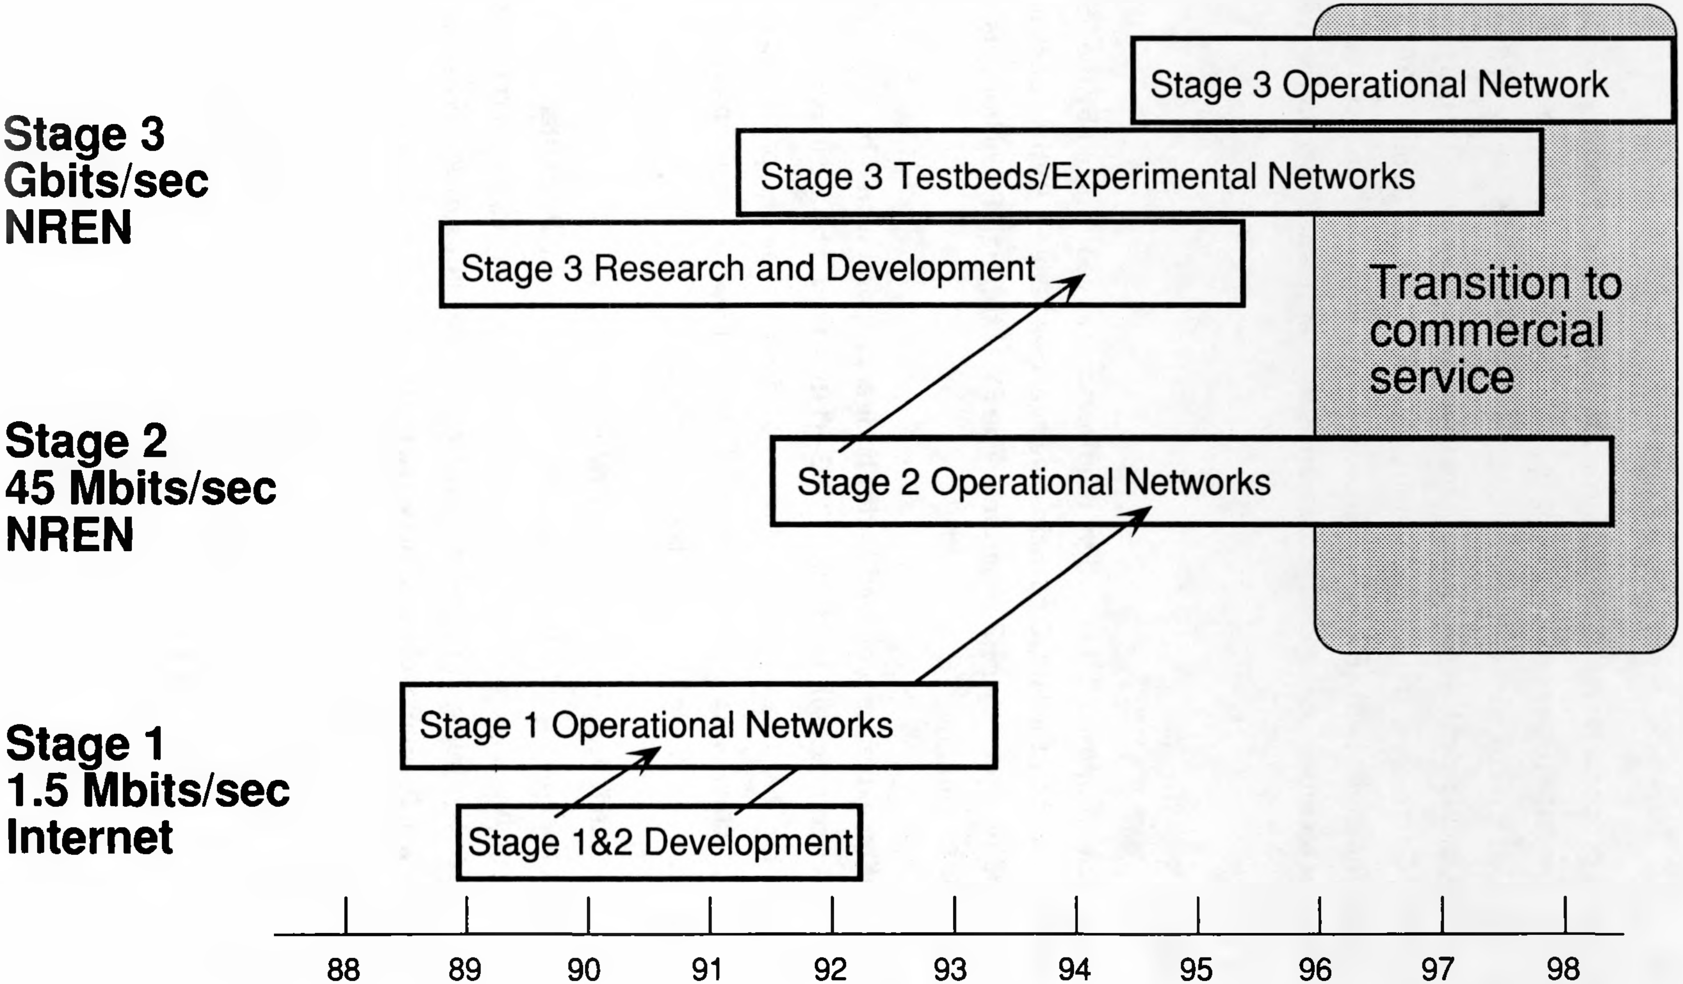

Figure 7-3. NREN implementation plan.

There now exists a proposal for a formalized structure to manage and coordinate the creation of NREN and the incorporation of existing U.S. research networks into NREN (see Figure 7-4). This structure is somewhat analogous to the ESnet structure in that the NREN user group, including three DOE representatives, will serve a similar purpose to the ESnet Steering Committee, and the Federal Engineering Planning Group (FEPG) will provide the same functions as the ESnet Site Coordinating Committee. In fact, it is intended that some members of the ESnet committees will represent ESnet on the NREN committees.

7.4 ESnet Budget Requirements

Table 7-2 presents the proposed ESnet budget requirements as coordinated by the ESnet Steering Committee and by the HPCC participating agencies. The FY 91 base does not include partial funding from other ER programs. The FY 92 through FY 96 budgets include HQ's base of $2 million.

Major implementation expenses proposed in this budget proposal would be (1) upgrades to T3 circuit capabilities during FY 92–93 (this would be accompanied by higher capacity routing hardware), (2) introduction of visualization and packetized video conferencing technologies in FY 93–94, and (3) initial tests and deployment of gigabit technologies in FY 95–96 or as they become available.

The DOE gigabit research program is listed separately in Table 7-2. This work will concentrate on end-user and applications support. It will focus on distributed computing environments, remote experiment control and monitoring, and services for supporting high speed coupling of remote supercomputer and other massively parallel processors, such as windowing, tools, and network protocols.

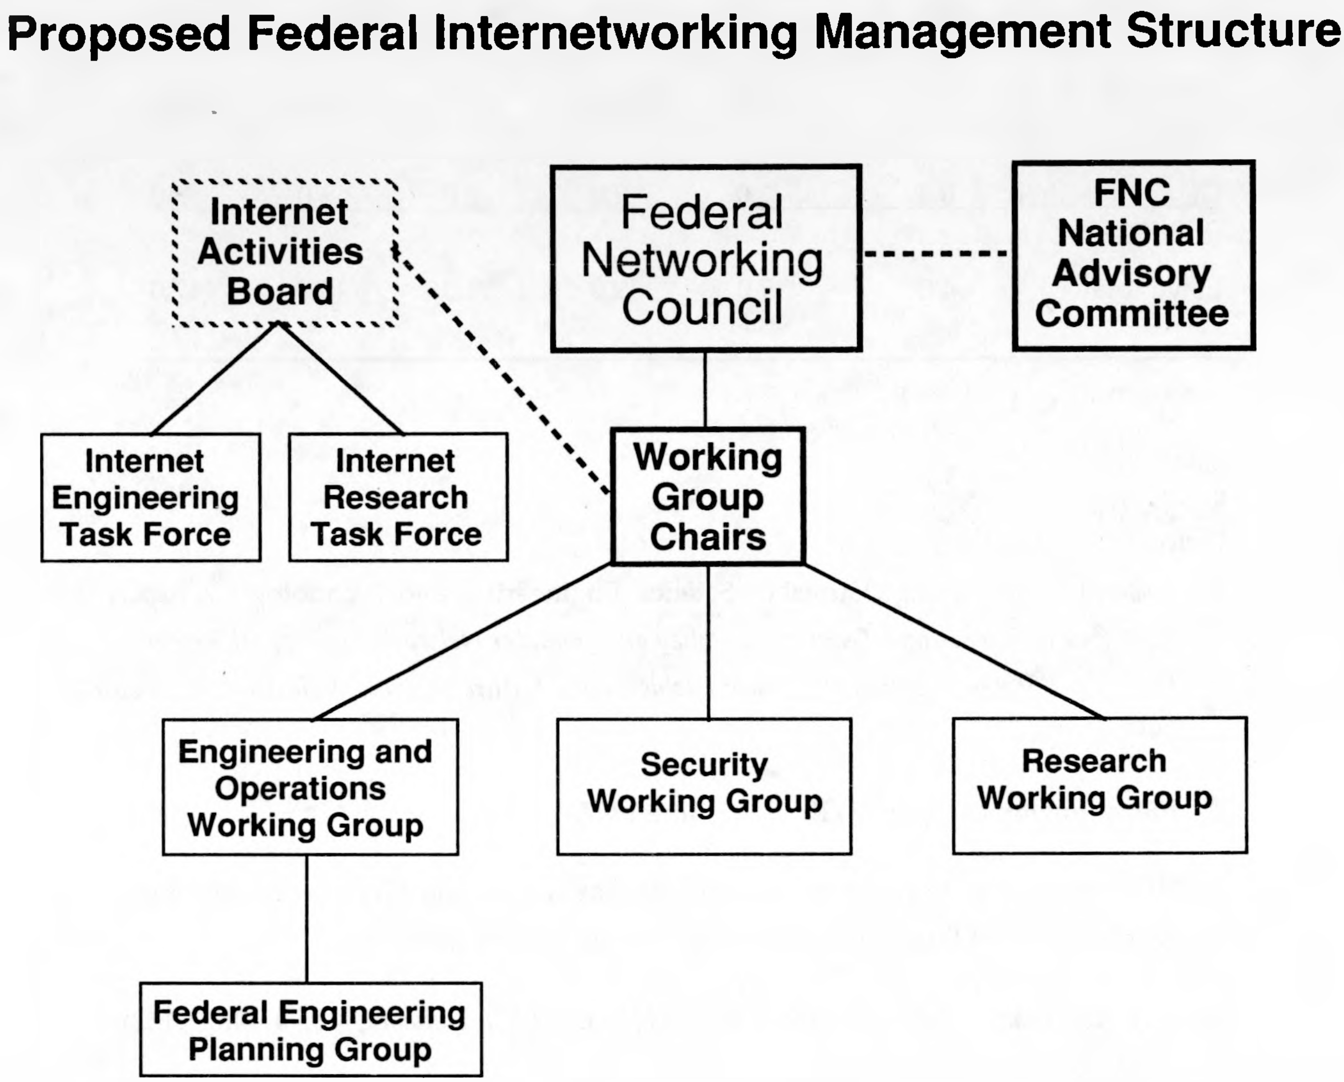

Figure 7-4. Proposed Federal internetworking management structure.

Table 7-2. Proposed ESnet budget requirements (M$).

Funding Component	FY 91 base	FY 92	FY 93	FY 94	FY 95	FY 96
Total operations	3.6	7.0	10.0	12.0	13.0	14.0
Gigabit research	0.4	2.0	2.0	2.0	3.0	3.0
ESnet total	4.0	9.0	12.0	14.0	16.0	17.0

References:

1. Federal Coordinating Council on Science, Engineering, and Technology, *A Report to the Office of Science and Technology Policy on Computer Networks to Support Research in the United States, A Study of Critical Problems and Future Options*, Volume I, November 1987.

2. *ESnet Program Plan*, DOE/ER-0341, June 1987.

3. VHS tape, dated 1/4/90 Version 1 by the National Center for Supercomputing Applications on Study of a Numerically Modeled Severe Storm.

4. W.P. Lidinsky, "Data Communications Needs," *IEEE Network*, Vol. 4, No. 2, March 1990, pp. 28–33.

Appendix A:
Definition of Terms and Acronyms

A

Appendix A : Definition of Terms and Acronyms

From Nuclear Physics Computer Networking, DOE/ER-0458T.

The following is a list of terms and acronyms likely to be encountered in reading this report.

ACRF	Advanced Computational Research Facility
AMS	Applied Mathematical Sciences
ANL	Argonne National Laboratory
ANSI	American National Standards Institute
APS/DNP	American Physical Society/Division of Nuclear Physics
ARPAnet	Advanced Research Projects Agency Network
BARRnet	San Francisco Bay Area Regional Research Network
BES	Basic Energy Sciences
BITnet	Because It's Time Network
BNL	Brookhaven National Laboratory
CAD	Computer-aided design
CAE	Computer-aided engineering
CCIRN	Coordinating Committee for Intercontinental Research Networks
CCITT	International Telegraph and Telephone Consultative Committee
CDC	Control Data Corporation
CEBAF	Continuous Electron Beam Accelerator Facility
CERFnet	California Education and Research Federation Network
CERN	European Center for Nuclear Research (Centre Europeene pour la Recherche Nucleaire)
CHAMMP	Computer Hardware Advanced Mathematical and Model Physics
CICnet	Committee on Institutional Cooperation Network
CIT	Compact Ignition Tokamak
CLNP	ConnectionLess Network Protocol
CPU	Central Processing Unit
CRU	Computer resource unit
CTSS	Cray Timesharing System
DARPA	Defense Advanced Research Project Agency
DCA	Defense Communication Agency
DCE	Data Communications Equipment
DDCMP	Digital Data Communications Message Protocol (DEC)

DDN	Defense Data Network
DDS	Digital Data Service
DEC	Digital Equipment Corporation
DECnet	Communications hardware and software available on DEC computers
DESY	German Electron Synchrotron Laboratory, Hamburg, Germany
DOD	Department of Defense
DOD-IP	DOD Internet Protocol
DOE	Department of Energy
DTE	Data Terminal Equipment
EARN	European Academic Research Network
ECMWF	European Centre for Medium-range Weather Forecasts
EDA	Engineering Design Activity
EDWG	ESnet DECnet Working Group
EML	Environmental Measurement Laboratory
ER	Energy Research
ES	Energy Sciences
ESCC	ESnet Site Coordinating Committee
ES/DECnet	Portion of ESnet running DECnet
ESF	Extended Superframe Format
ESMAN	ESnet Manager and Staff
ESnet	Energy Sciences Network
ESNWG	ESnet NREN Working Group
ESSC	ESnet Steering Committee
ESWG	ESnet Security Working Group
FCCSET	Federal Coordinating Council for Science, Engineering, and Technology
FDDI	Fiber Distributed Data Interface
FE	Fusion Energy
FEPG	Federal Engineering Planning Group
FIPS	Federal Information Processing Systems
FIX	Federal Interagency Exchange
FNAL	Fermi National Accelerator Laboratory
FNC	Federal Network Council
FOPG	FNC OSI Planning Group
FRICC	Federal Research Internet Coordinating Council

FSU	Florida State University
FTP	File Transfer Protocol
FTS	Federal Telephone System
GA	General Atomics
GCM	General Circulation Model
GOSIP	Government OSI Profile
GSA	General Services Administration
HDLC	High-level Data Link Control
HENP	High Energy and Nuclear Physics
HEP	High Energy Physics
HEPAP	High Energy Physics Advisory Panel
HEPnet	High Energy Physics Network
HER	Health and Environmental Research
HHS	Health and Human Services
HISS	Heavy Ion Spectrometer System (at LBL's Bevalac)
HPC	HEPnet Policy Committee
HPCC	High Performance Computing and Communications Program
HRC	HEPnet Review Committee
HTCC	HEPnet Technical Coordinating Committee
IAB	Internet Advisory Board
IBM	International Business Machines
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
Internet	Collection of packet-switched and broadcast networks
IP	Internet Protocol
ISO	International Standards Organization
ISODE	ISO development environment
ITER	International Thermonuclear Experimental Reactor
IUCF	Indiana University Cyclotron Facility
JET	Joint European Torus, Culham Laboratory, England
JNET	Commercial vendor for BITnet software
Kbps	Kilobits per second
KEK	Japanese High Energy Physics Laboratory, Tsukuba, Japan
L3	CERN LEP Experiment #3

LAN	Local Area Network
LAMPF	Los Alamos Meson Physics Facility
LANL	Los Alamos National Laboratory
LAPB	Link Access Procedure (Balanced)
LAT	Local Area Transport
LAVC	Local Area VAX Cluster
LBL	Lawrence Berkeley Laboratory
LEP	Large Electron-Positron Ring (CERN)
LEP3net	Network for CERN L3 experiment
LLNL	Lawrence Livermore National Laboratory
LREH	Laboratory for Radiological Biology and Environmental Health
Mbps	Megabits per second
Mb	Megabytes (8 bits/byte)
MFE	Magnetic fusion energy
MFECC	Magnetic Fusion Energy Computer Center
MFEnet	Magnetic Fusion Energy Network
MFEnet II	Proposed expansion of MFEnet
MIDnet	Midwestern States Network
MILnet	Military Network
MIPS	Million instructions per second
MIT	Massachusetts Institute of Technology
NAC	Network Access Committee
NASA	National Aeronautics and Space Administration
NEARnet	New England Academic and Research Network
NERSC	National Energy Research Supercomputer Center
NFS	Network file system
NIST	National Institute of Standards and Technology
NMFECC	National Magnetic Fusion Energy Computer Center
NMR	Nuclear Magnetic Resonance
NNC	National Network Council
NNDC	National Nuclear Data Center
NNT	National Network Testbed
NOCC	ESnet Network Operations Control Center
NORTHWESTnet	Northwestern States Network

NP	Nuclear Physics
NPNC	Nuclear Physics Network Coordinator
NPPCN	Nuclear Physics Panel on Computer Networking
NREN	National Research and Education Network
NSAC	DOE/NSF Nuclear Sciences Advisory Committee
NSF	National Science Foundation
NSFnet	National Science Foundation Network
NSN	NASA Science Network
NSP	Network Services Protocol
NYSERnet	New York State Educational Research Network
OER	Office of Energy Research
OHER	Office of Health and Environmental Research
OIW	NIST OSI Implementors Workshop
OMnet	A third-party Telenet-based mail/bulletin board service for Scientific users
ORELA	Oak Ridge Electron Linear Accelerator
ORNL	Oak Ridge National Laboratory
OSSC	Office of the Superconducting Super Collider, Germantown, MD
OSF	Open Systems Foundation
OSI	Open Systems Interconnect
OSI CLNP	Open Systems Interconnection ConnectionLess Network Protocol
OSTP	Office of Science and Technology Policy
OTI	Satellite link vendor
PACX	Packet exchange
PAD	Packet Assembler Disassembler
PC	Personal computer
PET	Positron Emission Tomography
POSIX	Portable Operating System Interface Exchange
PPPL	Princeton Plasma Physics Laboratory
PSI	Packet System Interface
QSPIRES	Remote access to SPIRES system (bibliographically oriented database at Stanford University)
RHIC	Relativistic Heavy Ion Collider (to be built at BNL)
RISC	Reduced instruction set computing
RPC	Remote procedure call

RS-232	Established definition of physical layer interface for interconnection of equipment
RSCS	Remote Spooling Communication System
RSX	DEC operating system for PDP computers
RT	DEC operating system for PDP computers
SCRI	Supercomputer Computations Research Institute
SCS	Scientific Computing Staff
SESQUInet	Texas Sesquicentennial Network
SLAC	Stanford Linear Accelerator Center
SMTP	Simple Mail Transfer Protocol
SNA	Systems Network Architecture (IBM)
SPAN	Space Physics Analysis Network
SPS	Super Proton Synchrotron (at CERN)
SREL	Savannah River Ecology Laboratory
SSC	Superconducting Super Collider
SSCL	Superconducting Super Collider Laboratory
SURAnet	Southern Universities Research Association Network
T1	1.544 Megabit-per-second circuit
T3	44.736 Megabit-per-second circuit
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TELNET	Remote login
TFTR	Tokamak Fusion Test Reactor
TISN	Japanese wide area net
TRIUMF	Canada's National Meson Research Facility
THEnet	Texas Higher Education Network
UNIX	A timesharing operating system developed by Bell Laboratories
USGS	United States Geological Service
UT	University of Texas
UTA	University of Texas at Austin
UUCP	Unix to Unix Copy Program
VAX	DEC computer system
VMS	DEC operating system for VAX computers
WAN	Wide Area Network

WESTnet	Southwestern States Network
WIDE	Japanese wide area net
WIN	German wide area net
X-11	A distributed windowing and graphics software package
X.25	Packet mode DTE to DCE interface
X.29	DTE/PAD communications protocol
X.500	Directory services
X.400	Electronic mail standard
XNS	Xerox Networking Systems

Appendix B:
ESnet Usage Guidelines and Procedures

B

Appendix B: ESnet Usage Guidelines and Procedures

Introduction

Scope

The purpose of this document is to convey guidelines and procedures regarding use of the Energy Sciences Network (ESnet). These guidelines are subordinate to all applicable U.S. government laws, as well as Department of Energy (DOE) orders.

ESnet users are generally DOE Energy Research (ER) supported researchers who rely on the availability of a production computer network. It is the responsibility of the ESnet Steering Committee (ESSC) and the ESnet implementers to ensure that the use of ESnet by any individual researcher is accomplished in a manner that does not unduly affect other network users. Any restriction of use of the network contained herein is intended to protect this ER resource for its intended use.

ESnet is operated by the National Energy Research Supercomputer Center (NERSC). ESnet management is guided by the ESnet Steering Committee, appointed by the DOE Scientific Computing Staff, with representatives from each of the Energy Research programs.

Questions regarding these guidelines and procedures may be directed toward the ESnet Network Manager, either electronically (info@es.net for Internet, 42158::info for DECnet), or using surface mail via the Energy Research Scientific Computing Staff, ER-7, GTN, Washington, D.C. 20545.

Purpose of ESnet

ESnet is a computer data communications network managed and funded by the Department of Energy, Office of Energy Research (DOE/OER), for the purpose of supporting multiple program, open scientific research. ESnet is intended to facilitate access to ER scientific facilities, provide needed information dissemination among scientific collaborators throughout all ER programs, and provide widespread access to existing supercomputer facilities.

The goals for the ESnet concept are:

- to enhance ER scientific collaborations nationally and internationally
- to provide equitable access to ER research facilities
- to avoid redundant computer network costs
- to support multivendor operations
- to accommodate the introduction of new technologies in a timely manner
- to interface to other agencies' networks
- to support distributed computing functions
- to provide significant contributions toward the establishment of a national research network.

Background

ER researchers have a strong history of wide-area-network usage. Many ER researchers frequently access the facilities of HEPnet, MFEnet, MILnet, NSFnet, and other university and government networks.

The MFEnet (Magnetic Fusion Energy Network) was created in 1976 to provide access to the National Magnetic Fusion Energy Computer Center (NMFECC, since renamed NERSC), located at Lawrence Livermore National Laboratory. The mission of NERSC and MFEnet since then has been extended to include all the DOE Energy Research programs. MFEnet, now redesigned to use ESnet, reaches all the national laboratories, a number of universities, the Supercomputer Computations Research Institute (SCRI) at Florida State University, and other domestic and foreign institutions.

MFEnet uses its own "NSP" protocols, which are implemented for VAX/VMS and the CTSS operating system used on the NERSC Crays. Remote login, file transport (export only), remote printing, and electronic mail are provided. There are mail gateways between MFEnet and Internet, BITnet, HEPnet, and internal networks at various large sites.

The ER High Energy Physics (HEP) research community has created the HEPnet (HEP Network) to support that community's networking needs. HEPnet refers to the entire HEP network infrastructure, including DECnet, BITnet, X.25, terminal access lines, dial-up lines, etc.

The major High Energy Physics laboratories (ANL, BNL, FNAL, LBL, SLAC) are interconnected with numerous universities and other institutions, as well as with a network of European HEP sites. The primary purpose of this network is to facilitate the geographically dispersed collaborations typical of HEP research projects. Services available include remote login, file transfer, and mail. Mail gateways exist between HEPnet and SPAN, BITnet, Internet, and MFEnet.

Description

The ESnet T1 (1.5 Megabit/second) backbone began initial operation in late 1989. ESnet provides access to all major ER laboratories and many other sites, including universities, private research institutions, and subcontractor and industrial collaborators, with projects supported by OER.

The ESnet backbone currently provides IP (Internet Protocol) and DECnet (Phase IV) packet switching services. These services support the higher-level DECnet protocols of HEPnet, and the NSP suite of protocols used by MFEnet, as well as the TCP/IP suite used by Internet. The goal is to provide increased connectivity and bandwidth for HEPnet, MFEnet, and users of the TCP/IP protocol suite in support of all the Energy Research programs.

Switching of X.25 packets was planned to be added in late 1990. Capability of supporting the International Standards Organization (ISO) Open Systems Interconnect (ISO) Connectionless Network Protocol (CLNP) is planned for mid-1991. Upgrades to bandwidth beyond T1 speeds are under active consideration.

Usage Guidelines

General Considerations

ESnet is not intended to compete with comparable commercial network services.

Users of ESnet shall not violate privacy or other applicable laws in making use of the network. ESnet shall not be used for advertising or other promotional purposes without the express permission of the OER Scientific Computing Staff.

Unauthorized use of ESnet is a Federal offense. Persons who break into government networks or use government resources without authorization may be prosecuted. Hosts or sites that knowingly or through negligence permit this type of usage will be disconnected from the network.

Classes of Usage

- *Class 1: traffic generated by usage in support of ER programs.* Network activity related to DOE/OER-supported programs constitutes authorized use of ESnet and is considered to be Class 1 usage.

- *Class 2: traffic generated by usage in support of either DOE non-ER programs or DOE authorized work for others.* Network activity related to DOE activities, including work for others, but not included in Class 1 usage, is considered to be Class 2 usage.

- *Class 3: traffic generated by all other usage.* ESnet will provide additional connectivity through interagency gateways and will participate in the evolution of the National Education and Research Network. These are examples of allowed Class 3 usage.

Note: Current network technology does not allow traffic to be identified by the above criteria. Nevertheless, at the policy level, the class of traffic is to be determined by actual usage, not by source or destination address.

Access

Authorization

Access to ESnet for Class 1 usage through a given site may be authorized by the ESnet Site Coordinating Committee member for that site. Access for usage of ESnet that will adversely affect the network may be denied, even though the application would constitute legitimate usage as defined heretofore.

Requests for access that will require new physical facilities or will significantly impact existing network facilities should be made to the ESnet Steering Committee member representing that program. The Steering Committee will prioritize all such requests and forward their recommendations to the OER Scientific Computing Staff (SCS).

Access to ESnet for Class 2 and 3 usage is authorized by the OER SCS.

Control

The ESnet access facilities at a user site shall be under the direct control of an Energy Research–funded organization of that site. The organization must also be willing to accept responsibility for coordinating the access requirements of all other ER programs at that site.

Security

Host Security

ESnet (by its nature) provides access for a much wider user base to connected resources. However, implementation or enhancement of end-node (e.g., host) security, which may be required as a result, is **not** considered to be a network-level issue, and responsibility must lie elsewhere.

The network will be responsible for providing reasonable tools and capabilities to allow the network operations staff to aid end nodes during periods of access intrusion.

Network Security

The network is recognized as a costly resource that must be protected from unauthorized usage.

Access to network control and operation functions must be protected to allow use only by authorized users.

End-node administrators are responsible for ensuring that network access and usage allowed through that end node meet the constraints and requirements of this document.

Data Security

Correct delivery and integrity of data submitted to the network for transmission is considered to be a network-level responsibility.

Additional data protection through means such as data encryption is **not** considered to be a network-level responsibility.

Data submitted for transmission by the network must be unclassified, non-sensitive information.

Physical Security

Sites hosting network communication equipment are required to provide necessary physical security.

Accounting

Collection of data necessary to account for usage of network resources is considered to be a network-level responsibility. Data collected will include source and destination identification, as well as traffic volume.

Traffic shall be identified as Class 1, 2, or 3.

Accounting for Class 2 traffic shall be reported to the OER/SCS.

Accounting for Class 3 traffic will be reported to the DOE representative to the Federal Network Council (FNC). Imbalances in reciprocal network services between agencies may thus be adjusted.

Naming and Addressing

Host naming and addressing shall be consistent with prevailing standards for the supported networks by hosts using ESnet.

Priority

Both access authorization and traffic handling shall be determined in a strict priority manner, such that Class 1 usage receives highest priority and Class 3 usage receives lowest priority. The intent is to allow access for and to carry Class 2 and 3 traffic only on a noninterference basis.

Foreign Traffic

Classification of Traffic

Foreign traffic resulting from a formal ER international agreement will be considered to be Class 1 or 2, as appropriate.

Other foreign traffic is defined to be Class 3 and will be supported contingent on the agreement of FNC.

Cost Sharing

Foreign-traffic network costs shall normally be divided at the half-circuit point of demarcation; i.e., the ESnet share of costs would include distribution costs within the U.S. plus the U.S. half-circuit costs.

Appendix C:
A Computer Networking Primer for Nuclear Physicists

C

Appendix C : A Computer Networking Primer for Nuclear Physicists

From *Nuclear Physics Computer Networking*, DOE/ER-0458T.

The Department of Energy's Nuclear Physics Panel on Computer Networking has prepared this summary in order to provide the working Nuclear Physics person with a general knowledge of the computer networking resources available to him or her, with some detailed information on how to access those resources.

Introduction

A computer network is a set of computers connected by physical transmission means and programmed to understand the transmissions (that is, the computers use common protocols). It is important to realize that, from the user's view, a specific "Network" is not a well-defined group of machines and unique connecting links. For example, new computers can be added, unannounced, to a network simply by dialing into a computer already connected. Furthermore, your computer can communicate with someone else's by means of numerous combinations of intermediary computers and devices that receive and relay the data you send and receive.

Nearby computers may be connected in local area networks (LANs), which in turn can connect to other computers (or networks) over wide area networks (WANs). Networks can also communicate with other networks that use different protocols via gateways. A gateway is a computer that is connected to two networks and can translate and retransmit messages between them. Later in this document, some very useful gateways are described that can give you convenient and reliable access to otherwise unavailable sites throughout the U.S. and the world.

There are three main tasks for which a network is most often used: electronic mail, file transfer, and remote login. Mail is the most commonly utilized and most widely available network service. It allows a user to send a message to another user, who can read and respond at his leisure. File transfer capabilities are probably the second most common service. They allow exchanging files of differing data formats including text (that is, containing printable characters) and executable files. Remote login is the capability that allows a user to access a remote host over a network as if he were connected to a direct terminal. Remote command execution allows a user to execute a command (for example, in order to submit a job) on a remote host. The latter two desirable features are not available on all networks discussed here. Computer conferencing (bulletin boards) is another useful network function similar to electronic mail, but is not widely supported on the networks discussed here.

The physical means and protocol definitions currently used to connect nodes (i.e., computers) are both undergoing rapid change. These changes are principally in the form of faster communication speeds and more standardized protocols. Typical WAN physical connections include leased telephone lines at 9600 bits per second (bps), high speed dedicated lines (56 Kbps), and T1 links (1.544 Mbps), as well as satellite links. The "backbone" of a network consists of the transmission lines and devices that carry the heavy data traffic between the principal switching sites (to which regional connections are made). It is analogous to the interstate highway system.

The WAN protocol sets that are of most interest to us here are TCP/IP, DECnet, and ISO/IP. Later in this summary, we review a few of the most important commands of TCP/IP and DECnet.

TCP/IP (Transmission Control Protocol/Internet Protocol) is used by ARPAnet and NSFnet. It is the most widely used protocol suite since it is in the public domain and is therefore adopted by many manufacturers of microcomputer software as well as university network designers. Familiar commands are "ftp" (file transfer protocol) and "telnet" (remote login).

DECnet is a (licensed) product of Digital Equipment Corporation (DEC) that is run chiefly on VAX computers. Due to the existence of tens of thousands of VAXs, including hundreds in use by Nuclear Physics groups, this protocol (and its related devices) is probably the one most widely used by the community. Users can send mail as well as do remote login via the command "set host."

ISO/OSI (International Standards Organization/Open Systems Interconnect) is presently under development and will eventually serve as an international standard. DEC is developing DECnet into compliance with the ISO/OSI standards. The OSI standard is a seven-layer model that defines the functions and interrelationships required in a communications transmission. These definitions can be compared to a railroad system, as shown in Table C-1.

Table C-1. Seven layers of OSI model.

OSI Layer	Name	Railroad Analogue
7	Application	Delivers freight to buyers
6	Presentation	Loads freight on trucks
5	Session	Opens boxcars
4	Transport	Freight yard dispatcher
3	Network device control	The switches
2	Data link control	The signal lights
1	Physical connection	The tracks

Description of WANs Important to Nuclear Physics

The Nuclear Physics community has generally met its networking needs through utilizing the capabilities developed by neighboring High Energy Physics groups or services available through campus- and laboratory-wide computing facilities. The WANs with which the Nuclear Physics community is most involved are ARPAnet, BITnet, ESnet, HEPnet, MFEnet, NSFnet, and SPAN. A general description of these networks can be found in the main text of the source document for this appendix. (See Section II of that document.)

For a general user, the current situation can be described roughly as follows:

If you are logged on to a Digital Equipment Corporation computer, you may be connected to "DECnet" and have direct access (with a remote account) to most Nuclear Physics experimental sites throughout the world by using DECnet commands. BITnet is probably implemented through the VMS Mail System using JNET software from Joiner Associates. Access to the TCP/IP world may be provided with MULTInet software.

If you are running on a campus or laboratory non-DEC computer (e.g., CDC or Cray), you may be connected to NSFnet and have direct access to most supercomputers and mainframe computers at universities and laboratories throughout the U.S. by using TCP/IP commands. BITnet is likely to be available.

Later we describe a few ways in which DECnet and NSFnet users can use gateways to access computers on the other network as well as how to log on to ES/DECnet from MFEnet.

Costs

Generally, costs to the user are not dependent upon volume or use, although this may change. Leased line costs for tail circuits to major network nodes are usually paid from normal funding for the research. Universities and laboratories generally fund their connections to regional networks and backbone sites out of overhead costs and do not directly "back charge" users. Institutions are charged by BITnet administration in proportion to their "size." NSF supports the NSFnet backbone and provides some seed money to regional networks. However, the regional networks are separately administered and must obtain their own continuing funding from their member institutions.

Accessing and Addressing

A computer on the network can be accessed by its "NodeID": either a node name or its node number (address). The advantage of using the node name is that it is easier to remember and will stay the same even if the computer it refers to changes location. However, it might be that the tables that tell the system what node number is being referenced by the node name have not been updated to include the node name you are using. It is therefore advantageous to have both available.

BITnet

NODE NAMES

Node name can be up to 8 characters.

NODE NUMBERS

There are none.

COMMANDS

BITnet is often accessed through the site's local network mail system. Thus, the specific commands to send and receive messages and files vary.

DECnet

NODE NAMES

Up to 6 characters.

NODE NUMBERS

Area.Node 64 areas each with a possible 1024 nodes each.

COMMANDS (VMS Operating System)

SET HOST NodeID

COPY, DIR: Use NodeID as part of filename

TCP/IP

NODE NAMES

Domains: Domain3.Domain2.Domain1.DomainRoot

NODE NUMBERS

Four fields: Fwide.F2.F3.Flocal

Each field between 0 and 255. Thus 4096 million possible nodes.

COMMANDS

telnet NodeID

open

ftp

help

ls

get

put

bye

MFEnet

NODE NAMES

Three characters.

NODE NUMBERS

One number: 0-127. Thus there are 128 possible nodes.

COMMANDS

NETTY nam

NETOUT lfn

Using an Ultrix Gateway

Users on ES/DECnet and NSFnet can take advantage of gateway machines, which understand both DECnet and TCP/IP protocols, are connected to both networks, and transparently transmit communications between them. In general, to use this gateway, the user simply specifies its "NodeID" as an intermediate node in his transaction and appends an exclamation point (!) at the end of the reference. The following example uses a DEC microVAX running the Ultrix operating system located at CEBAF. Alternatively FNGATE at FNAL can be used (192.31.80.145 or 46.7).

The NodeID of the CEBAF gateway is:

ES/DECnet :	43.479	CVU001
NSFnet:	129.57.1.10	suragate.cebaf.gov

ES/DECnet to NSFnet

You enter:	SET HOST 44511:: (43x1024 + 479 = 44511)
Suragate responds:	login:
You enter:	"NSFnet NodeID"!

You can now log on to the NSFnet Computer

NSFnet to ES/DECnet

You enter:	telnet 129.57.1.10
Suragate responds:	login:
You enter:	"DECnet NodeID"!

You can now log on to the ES/DECnet computer.

Mail and file transfer (ftp) can also be sent through the gateway simply by specifying it as an intermediate node and putting a "!" at the end.

MFEnet to ES/DECnet (cannot go in reverse direction)

There is another gateway at Fermilab that can be used to communicate over ES/DECnet from MFEnet. Its NodeID is:

ES/DECnet : 42.350	FNMFE
MFEnet:	FNA
You enter:	NETTY FNA
FNA responds:	DECnet host?[FNA]
You enter:	"DECnet NodeID" (No "!")

You can now log on to the ES/DECnet computer.

Other gateways exist between these networks.

Information Resources

1. The Bulletin of the American Physical Society, Membership Directory, contains BITnet addresses for those members who have provided them to the APS.

2. The Los Alamos Meson Physics Facility (LAMPF) Users Office maintains a list of about 500 BITnet addresses.

3. On a VAX/VMS Computer running JNET software, you can receive a useful description of BITnet by executing:

SEND NETSERV@BITNIC "GET BITNET USERHELP"

On an IBM/VM computer, execute:

TELL NETSERV AT MARIST GET BITNET USERHELP

4. Hedrick, Charles L., "Introduction to the Internet Protocols." This 29 page document gives a tutorial style description of TCP/IP. It is available via anonymous ftp from topaz.rutgers.edu. The path name is pub/tcp-ip-docs/tcp-ip-intro.doc. That is, if you are connected to NSFnet you can do:

```
telnet topaz.rutgers.edu
    anonymous
    "any identification"
```

get pub/tcp-ip-docs/tcp-ip-intro.doc
quit

5. The Wollongong Group's "Internetworking: An Introduction."

6. Quarterman, John S., and Hoskins, Josiah C., "Notable Computer Networks," Communications of the ACM **29**, 10 (October 1986), 932–970.

7. Jennings, Dennis M., Landweber, Lawrence H., Fuchs, Ira H., Farber, David J., and Adrion, W. Richards, "Computer Networking for Scientists," Science **231** (28 February 1986), 943–950.

8. "ESnet Program Plan," U.S. DOE Office of Energy Research Report, June 1987 (DOE/ER-0341). Available from the National Technical Information Service, U.S. Department of Commerce, Springfield, Virginia 22161. Describes the program plan and additional background material for DOE ER networking.

9. Van Name, Mark L., and Catchings, Bill, "The LAN Road to OSI," Byte (July 1989), 148–152.

10. Krol, E., "The Hitchhiker's Guide to the Internet," miscellaneous University of Illinois Report, August 25, 1987.

11. EXOS 8043-02, "TCP/IP Network Software For VAX/VMS Systems." Reference manual describing DEC's ethernet software and hardware.

Appendix D: OSI Model

D

Appendix D: OSI Model

From *Nuclear Physics Computer Networking*, DOE/ER-0458T.

The OSI (Open Systems Interconnection) network communications model consists of seven layers of protocols in the stack (Physical, Data, Network, Transport, Session, Presentation, and Application). Many industry groups and manufacturers' groups have been formed to facilitate interoperable, multivendor, computer-based information and communications systems. Interoperability is achieved by the use of a consistent set of worldwide agreed-upon standards developed by standards groups of the International Standards Organization (ISO) and adoption of compatible "profiles." Many international-standards groups are involved; these include the American National Standards Institute (ANSI) and the International Telegraph and Telephone Consultative Committee (CCITT), among others. The ISO standards have been completed through the Network layer and through about half of the Transport layer.

At the present time, in the absence of the full implementation of the complete ISO/OSI protocols throughout the world and with the plethora of noninteroperating network communications protocols, the process of internetworking is cumbersome at best and often impossible. The "gateway" translation mechanism poses problems and is really not applicable at the highest layers of a protocol stack because of the processing burden on the translating host. Also, an individual host speaking many different protocols to achieve connectivity with disparate remote hosts is not a very satisfactory nor pleasing solution to internetworking. To achieve the maximum benefit of expensive wide-area network connectivity, it is incumbent upon network users and developers to make the process as efficient as possible. The use of ISO standards will help mitigate these problems.

Nearly all vendors and users of computer software and hardware agree that true interoperability between disparate platforms, and in computer networking in particular, is a desirable goal, and that adherence to international agreed-upon standards, i.e., ISO/OSI, is an absolute must. Digital Equipment Corporation is no exception and, even now, provides products that conform. True compliance with ISO/OSI will be provided in the DECnet protocol suite with the release of DECnet/OSI, i.e., Phase-V DECnet, well before the rigidity of the FIPS and GOSIP standards become fully mandatory. A procedure to migrate ES/DECnet to ISO/OSI is available, and close cooperation with Digital Equipment Corporation has been achieved to provide for the orderly migration.

The following describes network layers in the OSI model.

Application layer: provides access to the OSI environment for users and also provides distributed information services; this layer handles applications such as file transfer or mail.

Presentation layer: provides independence to the application processes from differences in data representation (syntax); its purpose is to resolve differences in format and data representation.

Session layer: provides the control structure for communication between applications; establishes, manages, and terminates connections (sessions) between cooperating applications.

Transport layer: provides reliable, transparent transfer of data between end points; provides end-to-end error recovery and flow control.

Network layer: provides upper layers with independence from the data transmission and switching technologies used to connect systems; responsible for establishing, maintaining, and terminating connections. Internet Protocol (IP) occurs at this layer and is responsible for internetwork routing and delivery; it is sometimes referred to as "layer 3.5."

Data link layer: provides for the reliable transfer of information across the physical link; sends blocks of data (frames) with the necessary synchronization, error control, and flow control.

Physical layer: concerned with transmission of unstructured bit stream over physical medium; deals with the mechanical, electrical, functional, and procedural characteristics to access the physical medium.

Table D-1 (taken from *Data and Computer Communications*, William Stallings, Macmillan, 1988, p. 405) provides one way to compare the ISO model and TCP/IP (common networking protocol developed to handle internetworking, when two communicating computer systems do not connect to the same network). The ISO model is labeled OSI, and the TCP/IP model is labeled DOD.

Table D-1. Comparison of OSI and TCP/IP models.

OSI		DOD	
7	Applications	7	Process/ application
6	Presentations	6	
5	Session	5	Host-Host
4	Transport	4	
3	Network	3	Internet
2	Data link	2	Network access
1	Physical	1	

Appendix E:
Enabling the Next Generation of Scientific and Medical Imaging: The Essential Role of High Speed Computing and Communications

E

Appendix E: Enabling the Next Generation of Scientific and Medical Imaging: The Essential Role of High-Speed Computing and Communications

William Johnston, Imaging Technologies Group
Information and Computing Sciences Division, Lawrence Berkeley Laboratory

This paper will illustrate the central role that high-speed computing and high-speed networks will play in enabling the next generation of scientific and medical imaging. An environment that will enable this imaging, together with a collection of problems and directions for solutions, is presented. An example of a scientific/medical imaging application that is currently not possible due only to limitations in network bandwidth and computer network interface technology is also presented. The application is drawn from the Research Medicine and Radiation Biophysics Division of the Lawrence Berkeley Laboratory. This group has long been a world leader in the development of advanced medical imaging techniques [such as nuclear magnetic resonance (NMR) imaging, positron emission tomography (PET) imaging, bio-electromagnetic imaging, etc.], and in the use of high-energy particle-beam radiation therapy. The proposed computing environment, and the mechanisms for achieving it, are the result of insights gained from many years of work in networking, distributed computing, and imaging by the staff of LBL's Information and Computing Sciences and Electronics Engineering Divisions.

There is a class of scientific and research medicine imaging problems that is currently done poorly or not at all due to limitations in computer networking, computational capacity, and the scientific software environment. This research imaging environment is characterized by three elements that are typically geographically dispersed. These elements are:

(1) the imaging device (for instance, the 600-crystal PET scanner, or the LBL Advanced Light Source synchrotron ring), and their associated control systems;

(2) a supercomputer; and

(3) a high-performance graphics workstation.

Figure E-1 depicts this environment. The architecture of this system is described below and illustrated in Figure E-2.

(1) The source of the image data is a large, unique facility at a scientific/medical research institution. The data generated for a single "unit of interest" (either in time, or space, or both) are large, typically of the order of 1000 megabytes.

(2) The computation required to convert the data collected by the imaging device to viewable images (e.g., NMR reconstruction), or to do the simulations that are needed to otherwise interpret or utilize the data, requires a supercomputer. The supercomputer is needed for its large memory (gigabytes), or its high processor speed [hundreds of megaflops (millions of floating point operations per second)], and usually both.

(3) The research staff must have access to an analysis and visualization workstation in order to interpret the image data. The visualization is used to understand the highly complex imagery, to guide the operation of the imaging device, and directly for medical study, diagnosis, and treatment planning. This workstation must provide interactive visualization of the results of the reconstruction and simulation done on the supercomputer, and provide for interactive feedback to the imaging device and the object being imaged (a patient, in the case of a medical physics facility).

(4) A consistent, versatile, distributed, and rich suite of software is necessary to couple the three elements of the high-speed imaging environment in order to do the imaging, the image reconstruction, the image analysis and simulation, and imaging device control, and to allow for algorithm development. This last is a very important step. The imaging devices in question are research facilities, and the data always require substantial transformation before they are meaningful for human observation. The ability to easily and quickly prototype and modify software algorithms for this image reconstruction is crucial.

The current environment lacks all of the above elements. The sources of data are either isolated from the supercomputers or are communicated over low-speed networks; the supercomputers are at best connected to the graphics and analysis workstations via low- to medium-speed networks; adequate graphics workstations are just barely available; and the current software environment is not well enough integrated to permit the

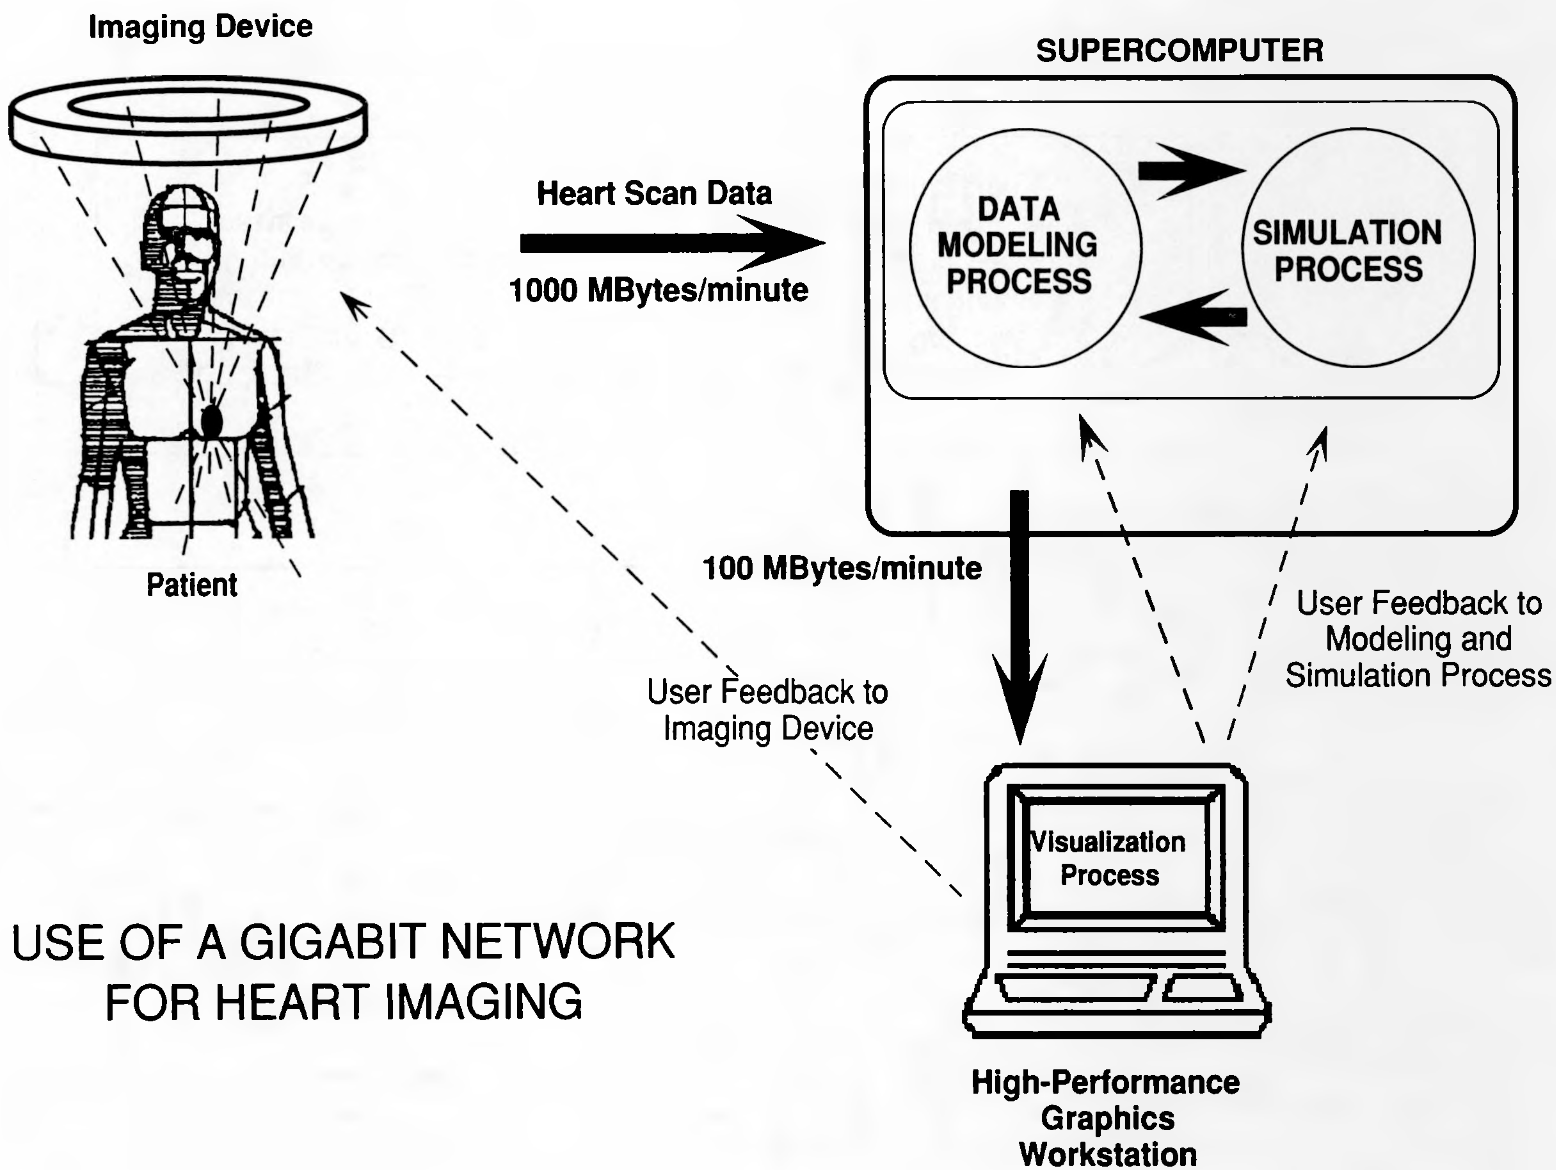

Figure E-1. Use of a gigabit network for heart imaging

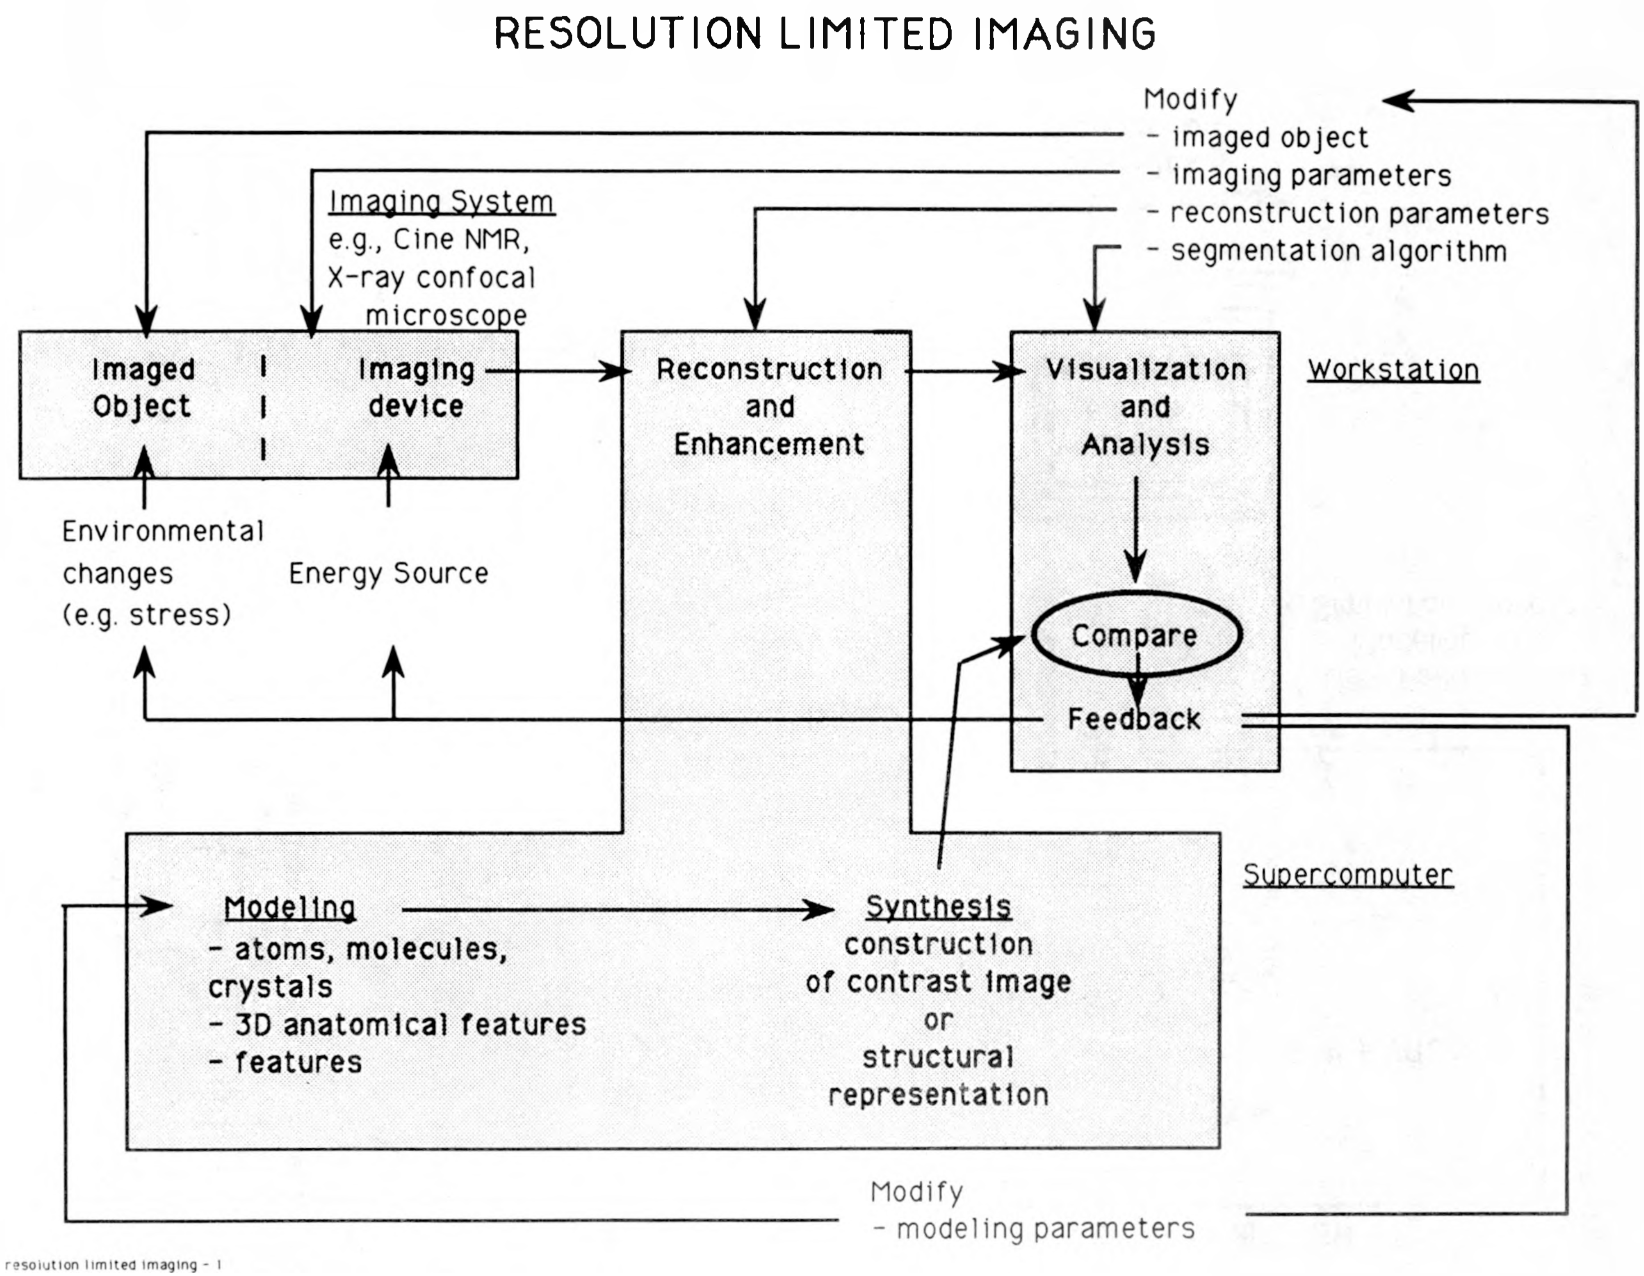

Figure E-2. Elaboration of imaging scenario shown in Figure E-1.

flexible development and experimentation needed for the various software algorithms that are at the heart of the imaging and visualization.

The solution to the first three of these problems involves eliminating a set of bottlenecks. As each bottleneck is relieved, another will emerge as the next problem to solve. Most of the issues are understood at this point (or, at least, are currently the topic of active research efforts), and the solutions are, for the most part, computer, electronics, or software engineering problems. There are no conceptual barriers to implementing the architecture described above.

The solution of each of these problems will lead to an immediate increase in the ability of the computing environment to support new scientific and medical physics imaging techniques. The solution of all of these problems will enable whole areas of scientific and medical investigation to develop.

Adequate network connectivity and bandwidth must be provided between the three geographically dispersed components. This entails not only network media and media access techniques, but also network management, high-speed protocols and routing algorithms, and high-speed network gateways. [1]

When the network bandwidth is adequate, the next bottleneck will be the computer/network interface. Interfaces and interprocess communication mechanisms must be provided that will permit computer systems to utilize high-speed data paths without so much overhead that all of the CPU resources are consumed just getting data from the network to computer memory to the user process, and back. [2]

Once the data paths operate at high speed, the next problem will be that of providing faster processing through both new and improved computer architecture and software algorithms (through the use of both coarse and fine grained parallelism, for example).

Finally, when all of the components are in place, we must have a software environment that will permit the integration of these components into a federation that will provide the distributed control, the data management, the modeling and simulation, and the visualization and analysis that are needed to solve the problem. [3]

Many examples could be provided to motivate the solution of the problems given above: holographic X-ray imaging of *in situ*, fragile biological specimens at LBL's Advanced

Light Source,[4] high-resolution confocal microscopy, patient monitoring and beam adjustment during high-energy particle-beam radiation oncolysis, multimodal imaging using combined bio-electromagnetic and NMR techniques for cardiac studies, real time NMR fluorescence studies of the heart and brain, and so forth. I have picked one example to elaborate on. This is a good example because it illustrates the real-time aspects of a potential medical physics application, it is easily quantified, and it could be done today, with no changes in the currently available medical imaging devices, if the high-speed imaging environment described above were in place.

In interactive angiography the scenario is to inject dyes distinguishable by NMR into a patient's blood vessels, and then to do the imaging and image reconstruction at a rate that allows real time studies: heart valve function studies, stressed heart studies, brain blood supply studies, and so forth. In other words, the function and behavior of the heart and its components could be studied directly under different environmental conditions. When this can be done in real time, or near real time, then the effects of regimes of exercise, drug treatment, or other forms of environmental stress can be observed directly, and the results incorporated into new forms of treatment, our knowledge of heart physiology, new diagnostic methods, etc.

To image a beating human heart with sufficient resolution to do heart valve studies, for example, one has to deal with a data transfer and processing rate that is dictated by the fact that the heart has a maximum excursion of about 1 cm, and a period of about one second. At thirty frames (NMR slices) per second, the heart could be imaged at a resolution of 256 x 256 x 6, once every 1/5 second. These conditions imply a data rate of about 500 megabytes/minute, or about 70 megabits/second. These data must be processed by (essentially) reverse Fourier transformation, which implies a computational speed of about 150 megaflops/second. The processed data (a three-dimensional image) must then be segmented and converted to a surface representation for graphical display. This implies about 100 megabytes/minute of graphics data that has to be sent from the supercomputer to a graphics workstation, there to be rendered, displayed, and interactively manipulated so that the research physician can direct the next phase of the experiment or diagnosis.

Research and engineering efforts are currently going on in the various computer hardware and software areas that can bring about the environment necessary to do this next generation of medical physics research. In order to validate this work, and to expedite and foster the computing environment necessary for real time medical imaging, there needs to be a cooperative effort to bring together the computing and communications components.

Of the many scientific imaging scenarios that require the described architecture, medical research can provide a clear and immediate focus to drive and justify the computing development, and demonstrate the utility and importance of an integrated, distributed, high-speed computing environment. High-speed network testbeds and associated prototype applications are the first step in the process of enabling this technology.

References:

[1] See, for example, "The Next Generation of MAN (Metropolitan Area Network)," A. Albanese, *The Bellcore Exchange*, September/October 1988; "Parallel Media Access Controller For Packet Communications at Gb/s Rates," Ippoliti and Albanese, *Proc. ICC/SUPERCOMM '90*; "An Analysis of TCP Processing Overhead," D. Clark, V. Jacobson, J. Romkey, and H. Salwen, *IEEE Communications Magazine*, June 1989.

[2] See, for example, "NETBLT: A High Throughput Transport Protocol," D. Clark, M. Lambert, and L. Zhang, *Frontiers in Computer Communications Technology: Proc. of the ACM-SIGCOMM '87*.

[3] See, for example, "The Software Bus: A Vision for Scientific Software Development," D. Hall, W. Greiman, W. Johnston, A. Merola, and S. Loken, *Computer Physics Communication*, 1989.

[4] See, for example, "The ALS Life Sciences Center," Lawrence Berkeley Laboratory PUB-5234, 1989.

Appendix F: Electronic Mail Addresses of ESnet Steering Committee Members and Selected Scientific Computing Staff

Mr. Robert J. Aiken
AikenRJ@es.net
(301) 353-5800; FTS 233-5800

Dr. Raymond A. Bair
d3e353@pnlg.pnl.gov
(509) 375-6807

Dr. George Brandenburg
brandenburg@huhepl.hepnet
(617) 495-2824; FTS 830-2824

Mr. John Cavallini,
cavallini@nersc.gov
(301) 353-5800; FTS 233-5800

Mr. Greg Chartrand
greg@sscvx1.hepnet
(214) 708-5090

Dr. James Davenport
davenport@bnl.mfenet
(516) 282-3789

Dr. Martin Greenwald
g@alcvax.pfcmit.edu
(617) 253-6053

Dr. Martha Hoehn
hoehn@lampf.lanl.gov
(505) 665-0637

Dr. Jean-Noel Leboeuf,
leboeuf@orn.mfenet
(615) 574-1127

Mr. James Leighton
jfl@nersc.gov
(415) 422-4025; FTS 532-4025

Mr. Sandy Merola (Chair)
axmerola@lbl.gov
(415) 486-7440; FTS 451-7440

Dr. Larry Price
lep@anlhep.hepnet
(708) 972-6295

Dr. Roy Whitney
whitney@cebaf2.cebaf.gov
(804) 249-7536

* U.S. G.P.O.:1991-281-714:40041

Appendix F:
Electronic Mail Addresses of ESnet Steering Committee Members and Selected Scientific Computing Staff

F

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

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, $300

ER-7