ESnet Program Plan 1994

U.S. Department of Energy
Washington, D.C.
Office of Energy Research
Office of Scientific Computing and
ESnet Steering Committee
November 1994

MASTER *ds*
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

For information concerning access to this document, see World Wide
Web universal resource locator (URL):
http://wwwosc.er.doe.gov/er30docs/Program_Plan_94.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Memorandum

Date: November 1, 1994

To: Martha Krebs, Director, Office of Energy Research
Dave Nelson, Office of Energy Research
John Cavallini, Acting Associate Director for Scientific
Computing, Office of Energy Research
George Seweryniak, ESnet Program Manager, Office of
Scientific Computing

On behalf of the Energy Sciences Network (ESnet) Steering Committee, I have the pleasure of submitting to you our ESnet Program Plan. This committee has been charged by the Office of Scientific Computing to represent and champion the computer networking and related services needs of the ESnet user community. It is on behalf of the programmatic needs of these Energy Research and other DOE Principal Investigators that this plan was created and submitted.

This Program Plan is distinguished by contributions from Principal Investigators throughout the ESnet community, and the preparation of the plan itself is a model of the collaboration that pervades this community. I know of no other program within the Federal government that is so dedicated to and comprehensive in its representation of programs, sites, and scientists.

This document will provide you with an overview of accomplishments, an updated vision, a motivation for continued growth, a changing ESnet paradigm, a few recommendations for your consideration, and—we on the committee would hope—your action.

As you peruse scientific journals, thumb through typical newsstand magazines, or converse with our nongovernment stakeholders (corporate executives and youths alike), you will find that computer networking is pervading our society, as well it should. It is also profoundly changing the way in which we work and collaborate. As you consider the scientific and business requirements of the Department of Energy, this plan demonstrates and codifies the necessity and priority of ESnet as a critical component of the Energy Research programs and infrastructure.

The ESnet Steering Committee and its associates look forward to the opportunity to provide any needed clarification, and we are available to help you accelerate the implementation of ESnet objectives. We are taking this opportunity to extend our appreciation for both the leadership and management of the Office of Scientific Computing in their acceleration of the Department's mission through the ESnet program.

On behalf of the ESnet Steering Committee,

Sandy Merola, Chairman

The Energy Sciences Network Steering Committee (ESSC) 1994

Raymond A. Bair, Basic Energy Sciences
Environmental Molecular Sciences Laboratory
Pacific Northwest Laboratory
Richland, WA
ra_bair@pnl.gov

Greg Chartrand, ESSC Executive Director
Computer and Network Technology Department
Pacific Northwest Laboratory
Richland, WA
greg@pnl.gov

James W. Davenport, Basic Energy Sciences
Department of Physics
Brookhaven National Laboratory
Upton, NY
daven@theory.phy.bnl.gov

John Drake, Health and Environmental Research
Mathematical Sciences Section
Oak Ridge National Laboratory
Oak Ridge, TN
bbd@ornl.gov

Martin Greenwald, Fusion Energy
Plasma Fusion Center
Massachusetts Institute of Technology
Cambridge, MA
g@cm2d2.pfc.mit.edu

Jean-Noel Leboeuf, Fusion Energy
Fusion Energy Division
Oak Ridge National Laboratory
Oak Ridge, TN
leboeuf@fedc04.fed.ornl.gov

Jim Leighton, Manager
Networking and Distributed Computing
National Energy Research Supercomputer Center
Lawrence Livermore National Laboratory
Livermore, CA
jfl@nersc.gov

Sandy Merola, Scientific Computing
ESSC Chairperson
Information and Computing Sciences Division
Lawrence Berkeley Laboratory
Berkeley, CA
merola@lbl.gov

Larry Price, High Energy Physics
High Energy Physics Division
Argonne National Laboratory
Argonne, IL
price@sg16.hep.anl.gov

Tom Rowlett, Human Resources
U.S. Department of Energy
Washington, D.C.
u7514tr@vm1.hqadmin.doe.gov

Walter Toki, High Energy Physics
Department of Physics
Colorado State University
Ft. Collins, CO
toki@lamar.colostate.edu

Robert Weiss, Health and Environmental
Research
Molecular Biology Institute
University of California at Los Angeles
Los Angeles, CA
weiss@uclaue.mbi.ucla.edu

Roy Whitney
ESnet Site Coordinating Committee Liaison
CEBAF Computer Center
Newport News, VA
whitney@asound.cebaf.gov

W.A. Zajc, Nuclear Physics
Nevis Laboratories
Columbia University
Irvington, New York
zajc@nevis.nevis.columbia.edu

Acknowledgements

This document was produced at the request of the Office of Scientific Computing, Office of Energy Research, of the U.S. Department of Energy. Editorial, graphic, and production services were provided by the Technical and Electronic Information Department (TEID) of the Information and Computing Sciences Division at Lawrence Berkeley Laboratory.

ESSC executive director Greg Chartrand acted as editor in chief for this project. Ron Fielder of LBL's TEID was responsible for rewriting and copy editing.

Contributors

ESSC Contributors

Raymond A. Bair
Greg Chartrand
John Drake
Martin Greenwald
Jim Leighton
Sandy Merola
Larry Price
Tom Rowlett
Robert Weiss
William A. Zajc

Steven Davis
Princeton University Particle Physics Laboratory
Princeton, NJ

Donn Davy
Human Genome Center
Lawrence Berkeley Laboratory

John Fitzgerald
National Energy Research Supercomputer Center
Lawrence Livermore National Laboratory

Tom Glanzman
Stanford Linear Accelerator Center
Stanford, CA

Other Contributors

Bob Aiken
Office of Scientific Computing
Germantown, MD

John Cavallini
Office of Scientific Computing
Germantown, MD

Les Cottrell
Stanford Linear Accelerator Center
Stanford, CA

William E. Johnston
Information and Computing Sciences Division
Lawrence Berkeley Laboratory

Stu Loken
Information and Computing Sciences Division
Lawrence Berkeley Laboratory

George Seweryniak
Office of Scientific Computing
Germantown, MD

This document incorporated the following electronic graphics files with the permission of the University of California, Lawrence Livermore National Laboratory, and the U.S. Department of Energy:

Multicast rgb.tiff
network services rgb.tiff

Contents

Executive Summary	ix
Introduction	1-1
ESnet History	2-1
ESnet Processes	3-1
Networking Committees.....	4-1
Collaborations and Liaisons	5-1
ESnet Today	6-1
High Energy Physics.....	7-1
Nuclear Physics	8-1
Fusion Energy.....	9-1
Basic Energy Sciences	10-1
Scientific Computing.....	11-1
Health and Environmental Research.....	12-1
Human Resources	13-1
Education	14-1
Toward the Virtual Workplace: New Modes of Scientific Collaboration	15-1
ESnet Tomorrow	16-1
Appendix: Acronyms and Abbreviations	A-1

Executive Summary

The Energy Sciences Network (ESnet) has distinguished itself as one of the world's most effective and progressive science-related computer networks. ESnet has played a major part in the development of the worldwide Internet, as it is now known, and ESnet continues to contribute to the future of networking through its participation in the High Performance Computing and Communications (HPCC) program, the National Information Infrastructure (NII) initiative, including its new Information Infrastructure Technology and Applications (IITA) component, and other related national and international enterprises. The advent of these initiatives brings ESnet and the entire Internet world to the threshold of major changes in technology, operations, and administration.

ESnet is a service-oriented production network chartered to support mission-oriented R&D. Its success consists not in the number of its links, its line speeds, or the number of protocols it supports but in its ability to provide services to its users. This success in providing standard as well as leading-edge services reflects an emphasis on user participation at every level of ESnet's structure. The members of the ESnet Steering Committee (ESSC), which defines requirements for ESnet and gives it general guidance, are associated with Energy Research (ER) programs and other

programs at the national laboratories and numerous academic research sites. Technical expertise for ESnet is contributed by the network experts from scientific laboratories and universities who participate in the ESnet Site Coordinating Committee (ESCC). Both committees cut across organizational and scientific boundaries, providing a rich mixture of users and experts who have developed a collaborative approach to sharing network facilities and services.

It is this user-driven, collaborative framework that sets ESnet apart from most other networks. This framework has positioned ESnet as the natural leader in promoting the cutting-edge technology that will be critical in the next phases of network evolution, as DOE contributes to the IITA initiative and the broader development of the NII.

Since its inception, ESnet has consisted largely of leased lines and hardware that connected Energy Research sites together in a complex network web. The current phase of network evolution will see these dedicated communications facilities replaced with a publicly shared communications infrastructure provided by private telecommunications vendors. Over the past three years, ESnet has been proactive in fostering its own evolution from a physical to a virtual network. Procurement complications have delayed the culmination of this transition, but once it is completed ESnet will have realized the vision of the emerging networking paradigm, wherein ownership of the network communications facilities is no longer required. ESnet's service-oriented philosophy has ensured its position of leadership during this period, as the importance of owning a dedicated physical network declines and the perceived value of higher-level services increases.

ESnet's evolution from a physical to a virtual network has facilitated the extension of its services to other DOE programs. Memoranda of Understanding governing provisory use of ESnet have already been established between ER and other program areas of DOE. These collaborations will lead to further shared use of infrastructure services, resources, and expertise.

This Program Plan is the product of the members of the ESnet Steering Committee and other ER principal investigators (PIs) who represent the broad user base of ESnet. These contributors wish to make the following recommendations:

1. We ask the DOE to incorporate an explicit NII component into the next version of the DOE Strategic Plan and to fully fund the ESnet activity as a critical component of the NII. In a large, distributed, heterogeneous organization like DOE, it is critical to recognize networking requirements and the need to disseminate information broadly as top-level strategic priorities.

2. We recommend that government leaders provide the forums needed to accelerate interagency collaborations in data communications, particularly with respect to increasing DOE's involvement in these initiatives. The Clinton/Gore administration is championing the benefits of computer networking to the nation's industrial, scientific, educational, and government sectors and to the population at large. The ESnet community can play a major role in advancing the administration's networking initiatives through the application of our leading-edge expertise and through our demand for advanced applications.

3. We recommend that DOE management assign adequate human resources to the Office of Scientific Computing (OSC) to support the efforts identified in this Program Plan. As ESnet's role has expanded from one of providing dedicated lines to one of creating and supporting distributed computing and information environments, the responsibilities of the ESnet Steering Committee have grown apace. As a result, some of our activities have come to require increased support from the OSC.

Introduction

This Program Plan characterizes ESnet with respect to the current and future needs of Energy Research programs for network infrastructure, services, and development. In doing so, this document articulates the vision and recommendations of the ESnet Steering Committee regarding ESnet's development and its support of computer networking facilities and associated user services. To afford the reader a perspective from which to evaluate the ever-increasing utility of networking to the Energy Research community, we have also provided a historical overview of Energy Research networking.

Networking has become an integral part of the work of DOE principal investigators, and this document is intended to assist the Office of Scientific Computing in ESnet program planning and management, including prioritization and funding. In particular, we identify the new directions that ESnet's development and implementation will take over the course of the next several years. Our basic goal is to ensure that the networking requirements of the respective scientific programs within Energy Research are addressed fairly. However, it should be kept in mind that while ESnet's basic mission is to support OER-funded research, other DOE offices are now participating in ESnet and sharing ESnet resources. These and other similarly evolving relationships are only in their infancy;

planning efforts for their further development are being initiated as this Program Plan goes to press.

This ESnet Program Plan is the third document of its kind. It was generated through the efforts of the ESnet Steering Committee and those of our associated researchers, managers, visionaries, and development and maintenance personnel. During the last few years, we have benefited from the vision and the implementation efforts of network experts both within the ER community and beyond. As the network has become increasingly transparent to users of more traditional computing tools, the advantages of networking have become highly visible and useful to those who create and reference essential scientific and administrative information and make it available to the ER community.

The proliferation of regional networks and additional network-related initiatives by other Federal agencies is changing the process by which we plan our own efforts to serve the DOE community. ESnet provides the Energy Research community with access to many other peer-level networks and to a multitude of other interconnected network facilities. ESnet's connectivity and relationship to these other networks and facilities are also described in this document.

Major Office of Energy Research programs are managed and coordinated by the Office of Basic Energy Sciences (OBES), the Office of High Energy and Nuclear Physics, the Office of Magnetic Fusion Energy, the Office of Scientific Computing, and the Office of Health and Environmental Research (OHER). Summaries of these programs are presented, along with their functional and technical requirements for wide-area networking. Changes in the ways these scientific programs use computing and information facilities have generally resulted in rapidly increasing networking needs. A major section of this Program Plan surveys current and future network utilization as projected by the Energy Research scientific community.

Forecasting networking demand is as much an art as a science. First of all, the measurement and analysis of current network utilization provides an important benchmark for estimating needed network growth. Second, many current end-user applications require bandwidth or connectivity that does not yet exist. Surveys of network users can help unearth such requirements. Third, advances in network-based applications and tools can generate less predictable accelerations in network utilization. For example, in this past year, we have benefited from the deployment of two new "killer" applications, videoconferencing and the network services related to Mosaic and the World Wide Web (WWW). These "killer" applications are so termed because their utility for and popularity with our scientific and information-based user community is generating increases in demand that will seriously stress the network unless bandwidth is increased quickly and dramatically.

The ESSC and its associates have long been active in generating and supporting Federal initiatives that pertain to networking. We are already benefiting from the early implementation of the Federal High Performance Computing and Communications Program and, in particular, its National Research and Educational Network component. We also hail the National Information Infrastructure as a vision that will further focus national attention on the critical role that networking must play in enhancing America's scientific and industrial competitiveness and in changing the way in which we work.

ESnet History

The impetus behind the formation of ESnet developed in the mid-1980s, when both the Fusion Energy (FE) and High Energy Physics (HEP) programs recognized the need for substantially improved computer network facilities. Until then, the Fusion Energy Community had been served by MFEnet, which was launched in 1976 as a result of the opening of a dedicated Fusion Energy supercomputer center at Lawrence Livermore National Laboratory (LLNL) in 1974. In order to make use of the new National Magnetic Fusion Energy Computer Center (NMFEC), FE researchers needed high-speed data links between their home sites and LLNL. This need was met by the initial MFEnet configuration, in which satellite links connected LLNL to a handful of key national laboratories and numerous tail circuits linked those labs to the other FE sites. By the mid-1980s, MFEnet had evolved from a medium for access to the NMFEC into a general-purpose network for Magnetic Fusion researchers.

HEP researchers had begun to use computer networking as soon as it became practical to do so, in the late 1970s. These first efforts involved microwave links between the Stanford Linear Accelerator Center (SLAC) and Lawrence Berkeley Laboratory (LBL). In the early 1980s a satellite link was established between SLAC and Argonne National Laboratory (ANL) to support a HEP experiment at SLAC. By the mid-1980s, the HEP program had developed an extensive network of leased lines (mostly operating at 9600 bps) that interconnected the main particle-accelerator

laboratories with numerous other sites. Until that time, ad hoc network management by volunteers from the HEP community had served the de facto HEPnet well. However, this system was expected to encounter serious difficulties managing the substantial upgrades that had become imperative as HEPnet utilization began to extend beyond the HEP community.

The early 1980s also saw other ER research programs joining established computer networks. Many university research groups began to use the electronic mail and file transfer facilities of BITnet or ARPAnet to communicate with their collaborators at the national laboratories. Other university groups found it necessary to lease direct connections to mainframe computers located at remote laboratories where those groups' research activities were concentrated.

The need for a more comprehensive approach to ER networking began to be appreciated in 1985. In that year, the HEP program established the HEPnet Technical Coordinating Committee to address HEPnet's intensifying management needs. In the same year, a subpanel of the DOE's High Energy Physics Advisory Panel (HEPAP) recommended the establishment of a formal HEPnet backbone to provide more effective and efficient networking for the HEP community.¹ This backbone was to consist of high-speed (56 kbps) trunk lines connecting the major HEP laboratories. This recommendation coincided in time with a proposal to upgrade the MFEnet.

Later in FY 1985, Dr. Alvin Trivelpiece, then Director of Energy Research, charged OER's Scientific Computing Staff (now the Office of Scientific Computing) with surveying computer networking requirements across all the ER programs and evaluating the status of existing network facilities. The results of this survey demonstrated that enhanced networking facilities were needed to improve access to unique ER scientific facilities, to facilitate the dissemination of information among scientific collaborators throughout all the ER programs, and to expand access to existing supercomputer facilities.

As a result of these results, Dr. Trivelpiece recommended that the MFEnet and HEPnet initiatives be combined into what would become ESnet, in order to optimize the efficiency and functionality of ER-wide networking. In a special presentation made in response to the survey's findings, the SCS staff set forth a number of more specific recommendations that became the foundations of ESnet. The staff recommended the formation of the Energy Sciences Network Steering Committee to represent the ER scientific community. The SCS staff also proposed an evolutionary model for the development of the new network and endorsed a phased approach to achieving long-term networking goals.

In October 1986, Dr. Trivelpiece approved a formal plan for the establishment of ESnet. The overall goal of the initiative was to create a single general-purpose scientific network for the ER community. The basic approach taken in organizing the new network was to combine the various ER programs' network activities by coordinating the applications-level requirements through the SCS staff. The ESnet Steering Committee was to be a source of guidance concerning these requirements and a source of general strategic oversight. The installation, coordination, and day-to-day operation of ESnet was to be the responsibility of the staff of the NMFEC, which had been renamed the National Energy Research Supercomputer Center (NERSC) to reflect its expanded role of providing supercomputer access and network services to a wider community.

The ESnet Steering Committee held its first meetings in late 1986, and ESnet began providing ER-wide networking services in January 1988. Initially, ESnet used time-division multiplexing across X.25 backbone lines, which operated at speeds of 56 kbps and 256 kbps. By 1989, ESnet had begun deployment of commercially supplied multiprotocol routers and T1 backbone lines, which provided speeds of 1.5 Mbps. This configuration became fully operational in early 1990, with 19 major OER-supported sites connected to the backbone. Although the performance level supported by the T1 backbone seemed quite respectable at first, by early 1991 it was becoming clear that a bandwidth upgrade would be required within a very few years, and planning for T3 (45 Mbps) capability began almost immediately.

During the late 1980s and the early 1990s, ESnet also began providing international connectivity in support of ER program activities. International connectivity was shared and coordinated with NASA, NSF (the National Science Foundation), and DARPA (the Defense Advanced Research Projects Agency), an approach that established a framework for future interagency cooperation. In the same time period, ESnet began connecting to regional NSF networks, thereby providing more ubiquitous network communications through which university researchers could utilize DOE/ER laboratories and facilities.

Today, as an integral part of the Internet, the Energy Sciences Network provides seamless, multiprotocol connectivity among a variety of scientific facilities and computing resources in support of collaborative research, both nationwide and internationally. ESnet also supports DOE-sponsored educational activities.

High-performance computing has now become a critical tool for scientific and engineering research. In many fields of research, computational science and engineering have become as important as the more traditional methods of theory and experiment. Progress and productivity in such fields depend on interactions between people and machines located at widely dispersed sites, interactions that can only

occur rapidly enough via the medium of high-performance computer networks. The ubiquity of networks has provided researchers with unexpected capabilities and unique opportunities for collaborations.

These benefits have only whetted the scientific community's appetite for still higher levels of networks performance to support wider network usage, the transmission of ever-greater volumes of information at faster rates, and the use of more sophisticated applications. The scientific community is also increasingly sensitive to the importance of protecting privacy and intellectual property. The mission of the Energy Sciences Network is to satisfy these needs as fully as possible for Department of Energy researchers.

Throughout the evolution of ESnet, its managers and the members of its committees have made significant contributions to the development of the worldwide Internet. Today the ESnet Steering Committee is keenly aware of the importance and scope of the proposed National Information Infrastructure and is formulating appropriate plans for OER's participation in NII development.

Notes

1. J. Ballam *et al.*, "Computing for Particle Physics," DOE/ER-0234, August 1985.

ESnet Processes

ESnet is a highly collaborative venture that involves the participation of technical experts, scientific users, and governmental liaison personnel. This broad spectrum of experts participates fully in the management, operation, and planning of ESnet. The collaborative quality of these processes ensures that user requirements are addressed as equably as possible, given ESnet's finite resources.

Because of its collaborative nature, ESnet must function within a framework of committees that facilitate user and technical participation in ESnet processes. The committees and the relationships they embody serve to focus the forces that are responsible for ESnet's success. These committees also provide the framework within which the ESnet community interacts, articulating its evolving vision of network support for scientific endeavor.

Planning Processes

Most of ESnet's planning processes are driven by the networking requirements defined by the ER programs that participate in ESnet. These requirements may originate from a variety of sources, but all are eventually reviewed by the ESnet Steering Committee. In general, new requirements are generated either by the introduction of new computer

applications or by the establishment of major new scientific projects or facilities.

Planning For New Applications

The march of technology spawns new computer and network applications that frequently challenge prevailing conceptions of network operation. In addition, new applications often generate increased network traffic loads that need to be evaluated. As new applications are introduced, they are assessed by ESnet personnel. If an application is deemed sufficiently demanding, ESnet initiates an evaluation/planning process, which may involve the participation of various ESnet individuals and committees. Many application-driven plans require the participation of technical experts who may reside at national laboratories or universities. The ability to form teams or task forces to plan for the implementation of new applications is the key to ensuring an effective response by ESnet.

Planning For New Scientific Projects and Collaborations

Major new projects and new scientific facilities usually generate additional network demands that require ESnet to initiate a planning process. Through its ESSC representative, a DOE program identifies its networking requirements for evaluation by the Steering Committee as a whole. Once these requirements are validated by the ESSC, responsibility for evaluation and implementation is passed on to ESnet management. If the new requirements have a major impact on the network, ESnet management will generally forward its implementation plans to the ESSC for review.

The Program Plan

ESnet reviews its requirements periodically through the process of revising its Program Plan, which identifies new applications, services, and operational parameters. Between Program Plan revisions, new services and requirements are usually documented in white papers, which are approved by the ESSC with the concurrence of the OSC and the ESnet management staff at NERSC.

Organizational Processes

The ESSC and the ESCC may collaborate to address an emergent situation that requires a cross-organizational response. Such a response

usually entails leadership by a member of either the ESSC or the ESCC. That committee member becomes that committee's designated liaison and attends the other committee's meetings to report on activity and to discuss issues.

Ad hoc subcommittees are formed within the ESSC and the ESCC to address specific issues and services. The chairpersons of the ESSC and the ESCC determine the need for such subcommittees and appoint their leaders. If there are support issues to be resolved, a representative of ESnet management will usually participate in the subcommittee.

Operational Processes

Operational processes are generally defined by the ESnet staff, with complex processes documented in advance and reviewed by a subcommittee or an interest group. There are special technology committees associated with several specific ESnet services, such as videoconferencing, directory services, and Open Systems Interconnection (OSI) implementation. Appropriate ESnet staff members participate in these working groups, which make it possible for service providers, end users, and facility management personnel from national labs and universities to work together very closely. This direct interaction optimizes the provision of services and expedites changes in operational processes.

Problem Resolution

Situations occasionally arise that require focused administrative and/or technical attention. Historically, ESnet has responded to these situations by forming task forces or term-limited subcommittees. If the issue is requirement driven or is a policy or administrative matter, the subcommittee may include members from the ESSC. If the issue is technical, the subcommittee is likely to include members of the ESCC. If necessary, membership can include members of both committees or noncommittee personnel. Generally, a task force will either produce a white paper or present its findings in person to the pertinent committee(s).

Performance Measurement

Through its Steering Committee, ESnet operations are reviewed periodically to assure that management, technology, and capabilities are adequate to meet the requirements detailed by the participating ER programs. As a result of both technical and administrative reviews, ESnet

has modified its plans to accommodate new services and utilize new technology. The most recent review has defined the basis on which ESnet will move forward into the NII era.

As a result of the February 1994 review, it has been determined that ESnet management will provide a yearly report of objectives and goals by which ESnet accomplishments can be measured. This report, along with the ESnet Annual Report produced by the ESCC, will document the progress and deployment of new technology, services, and capabilities.

ESnet also generates several performance measurements that are used for network management. These measurements are made publicly available on file servers. Alternatively, they may be electronically distributed on a periodic basis. These measurements are useful in determining whether network traffic patterns are changing. They also document the effects of new users, programs, and facilities.

Networking Committees

A successful service-oriented community must establish current and future user requirements and effectively satisfy those requirements within the operative funding and administrative constraints. In the case of ESnet, the processes of determining and satisfying requirements entail the most widespread collaborative efforts within the Department of Energy. Descriptions of current and future user requirements must be obtained from DOE-supported principal investigators who are located at national laboratories and universities across the U.S. and at other research sites throughout the world. The work of satisfying these requirements is concentrated at the National Energy Research Supercomputing Center at Lawrence Livermore National Laboratory, although this work also requires the efforts of network specialists at all major DOE sites and at other sites where important collaborators do their work.

The breadth of services provided by ESnet—from basic bandwidth to directory services to e-mail and videoconferencing support—requires a technically diverse support staff. The complex interactions of these technical experts are overseen by an effective, efficient committee structure whose hallmarks are cooperation and synergy.

Figure 4-1. Relationships among DOE networking committee

Figure 4-1 represents an ESnet-centered view of the many collaborating networking committees within DOE. It should be noted that several of the committees represented in this chart (SAC, CCIRDA, EXERSUG, ERSUG, and SCIE) are shown only to give the reader a sense of ESnet community interactions. These committees have major responsibilities that are only briefly explained in this document.

The Office of Scientific Computing (OSC)

ESnet is sponsored by the Office of Scientific Computing of DOE's Office of Energy Research. The OSC is therefore responsible for funding the network and overseeing its management. OSC's general charter is to improve DOE's ability to solve critical scientific and engineering problems by supporting research in the advanced mathematical, computer, and computational sciences and by fostering applications derived from the results of such research. As a significant component of this charter, OSC is responsible for managing the Energy Research supercomputer facilities and the Energy Sciences network.

The ESnet Steering Committee (ESSC)

The ESnet Steering Committee was originally formed in 1986 with representation from the ER program offices. As ESnet has extended its service to the EM (Environmental Waste and Restoration Management) and HR (Human Resources and Administration) communities, the ESSC

has expanded to ensure adequate breadth of program representation. The ESSC is charged to:

- Document, review, and prioritize network requirements for all ER programs
- Review the ESnet budget as presented by the Office of Scientific Computing, evaluating that budget with respect to the prioritized network requirements.
- Identify network requirements that require further research
- Establish performance objectives for ESnet
- Propose innovative techniques for enhancing ESnet's capabilities
- Advise the ESnet management personnel at NERSC

Members of the ESSC represent specific DOE program offices, and they are appointed by their Division Directors, at whose pleasure they serve. There are no firm guidelines as to length of service of an ESSC member. Experience has shown that the committee benefits equally from the consistency of purpose and direction provided by long-term members and the innovations fostered by new representatives. The OSC nominates an ESSC member to serve as committee chairperson, and the selection becomes final upon approval by a vote of the entire ESSC.

The ESSC has codified the criteria by which priorities can be set for major network improvements, including major new network connections as well as upgrades of existing facilities. These criteria include such factors as the importance of the facilities in question to DOE's mission and the ability of the requesting site to utilize and/or support the new or expanded network connectivity.

The ESSC has also formalized the process by which network requirements are prioritized and has documented these procedures in the minutes of its meetings.¹ This documentation defines the role of the OSC as the program office with oversight authority, the role of the ESSC as the representative of ESnet's users, and the role of ESnet management in implementation.

The ESnet Steering Committee continually addresses issues of significance for the future of DOE networking. Recently, these issues have included:

- Telecommuting
- Conference room and desktop videoconferencing services
- A significant upgrade of ESnet's bandwidth and services
- The development of an Acceptable Usage Policy for ESnet
- Increased coordination with SCIE and EXERSUG
- The need for DOE-wide coordination of distributed computing services

- An integrated approach to DOE's international connectivity
- The future effects of other networks on the ER community
- The contributions ESnet can make to the National Research and Education Network and the National Information Infrastructure, and the benefits of such participation
- The effects of related Federal activities, such as the Federal Networking Advisory Committee and the Federal Internetworking Requirements Panel
- The role of ESnet in education

The ESSC is also considering a number of strategic issues, such as:

- The development of performance metrics
- Increased connectivity to other network providers, both federal and commercial
- The importance of an improved public awareness strategy
- Network development as a key element in DOE's Strategic Plan
- Increased documentation of strategic and operating plans

The ESnet Site Coordinating Committee (ESCC)

The ESnet Site Coordinating Committee is composed of representatives from each of the major ESnet backbone sites. Established in 1987, the ESCC serves as:

- An advisory body to the ESnet Steering Committee, providing a forum for the consideration of a broad range of technical issues
- A forum for information interchange about ESnet-wide activities and plans and site-specific requirements and plans
- A forum for interactions with the ESnet manager and staff
- A forum for interactions with ER programs that use or would like to use ESnet facilities.

The ESCC is a standing committee whose members are appointed by the individual ESnet site organizations. Current membership represents 22 Energy Research sites. The ESCC chairperson is appointed by OSC from among the members of that committee, with the advice and consent of both the ESSC and the ESCC.

To carry out the functions listed above, the ESCC appoints various working groups and task forces as the need arises. A working group generally exists for an extended period of time to address issues within a general category (e.g., ESnet/DECnet issues or security issues). A task force, on the other hand, exists for only a short term to accomplish a narrow, well-defined goal (e.g., DECnet routing or TCP/IP routing). The

membership of working groups is subject to ESCC approval. Working groups choose their own chairperson, subject again to ESSC approval. Task forces may have a more flexible structure, depending on the issue, but their leadership and membership are subject to ESCC approval.

The Distributed Computing Coordinating Committee (DCCC)

Many of the activities associated with DOE programs have become global in scale, encompassing research, development, and construction around the world. A ubiquitous distributed computing environment has become a necessity to effectively manage and operate these global enterprises. The functioning of such an environment requires more than simply a robust physical network. It also demands that all layers of the network structure—from the physical layer to the application layer—be organized and managed properly. The DCCC was formed to address the problems associated with the higher levels of this structure.

To assure the necessary full range of implementation, three computer- and network-related committees—the ESSC, EXERSUG, and SCIE—have joined to charter the DCCC. Through the DCCC, these committees will actively work to facilitate the development of a functional distributed computing environment for the sciences associated with DOE research.

The Coordinating Committee for Informatics Research, Development, and Application (CCIRDA)

In addition to networking, the Office of Energy Research is responsible for providing the distributed technological infrastructure that underlies DOE's mission in computation and information delivery. The Coordinating Committee for Informatics Research, Development, and Application (CCIRDA) is charged with planning and coordinating efforts to satisfy informational, computational, and communications requirements throughout ER. In carrying out this mission, CCIRDA coordinates ER-wide Information Technology Infrastructure activities and assists in coordinating the plans, policies, and issues of Information Resource Management (IRM) with those of Management and Operating (M&O) Contractors within ER. These coordination responsibilities entail such tasks as helping to provide feedback and direction for DOE's Information Technology programmatic mission. These efforts are intended to support the development of an integrated long-range IRM science and technology plan and to contribute to the identification of issues and emerging technologies that will be critical to all DOE programmatic missions.

The ESSC chairperson serves as liaison with CCIRDA and attends the meetings of that committee. The members of CCIRDA have standing invitations to attend ESSC meetings.

ERSUG and EXERSUG

The Energy Research Supercomputer Users Group (ERSUG) consists of all investigators who use the supercomputing facilities provided by the Office of Energy Research. The mission of ERSUG is to promote the effective use of supercomputing facilities by sharing information about notable accomplishments, the capabilities and limitations of the supercomputer facilities, and new opportunities arising from hardware and software advances. ERSUG monitors the computational needs of the Energy Research community and communicates this information to its membership, to service providers, and to appropriate DOE officers. ERSUG also participates in developing multiprogram requirements for shared facilities. These requirements are coordinated with the end-to-end computing plans of individual programs. ERSUG also documents the progress afforded by the supercomputing resources and the future opportunities for research that require enhancements to these resources.

The organizational body for ERSUG is its executive committee (EXERSUG). Membership in EXERSUG is defined by programmatic constituency and by professional specialty. The Supercomputer Access Committee (SAC) determines the allocation of constituencies to be represented on EXERSUG. An EXERSUG member can be nominated by any ERSUG member from his or her own programmatic constituency, by the program office, or by an active member who wishes to be replaced. New nominees are elected or rejected by a vote of the current EXERSUG membership. Members usually serve for a period of three years or until a replacement is determined. EXERSUG may spawn subcommittees of appropriate experts and interested parties to address specific problems or issues.

The Scientific Computing Information Exchange (SCIE)

The Scientific Computing Information Exchange was begun in the early 1960s as a means of helping the directors of national lab computer centers resolve common problems by sharing information. Today SCIE remains a forum without a formal charter but with a keen interest in dealing with Information Resource Management issues. SCIE is currently taking the lead in addressing such issues as Information Technology Resource Planning and the privacy of electronic mail.

The Supercomputer Access Committee (SAC)

The Supercomputer Access Committee is composed of representatives of each of the ER Program Offices. Operating under guidelines provided by the OER, SAC is responsible for allocating the resources of the ER-funded supercomputer access centers to ER-funded investigators. SAC meetings provide a forum in which ER program managers can address issues related to the use of the supercomputer access centers, for example, assessing the computing capabilities that will be required to fulfill ER programmatic missions in the future.

Note

1. Minutes of ESSC meetings are available on the World Wide Web at the following universal resource locator (URL):
<http://www.es.net/hypertext/essc.html>

Collaborations and Liaisons

ESnet and its constituents participate in every conceivable aspect of networking. This activity involves working with many entities outside the DOE sphere. Such external activities involve the full spectrum of ESnet participants, from DOE and ESnet management personnel to ER researchers and computer professionals at national laboratories and universities. ESnet also provides forums for sharing the information gained from these liaisons. Access to this information is one of the major benefits to be derived from voluntary participation in ESnet. The major non-ESnet participants in these collaborations and liaisons are described below.

Committee/Board/Panel Participation by the Office of Scientific Computing

The Office of Scientific Computing participates in numerous federal and interagency initiatives and committees by providing technical and administrative support for these efforts, which include the following agencies and ad hoc groups:

- *The National Science & Technology Council (NSTC)* (formerly FCCSET). The NSTC is charged with developing a national R&D

strategy to coordinate federal R&D expenditures. Of the NSTC's nine subcommittees, the OSC participates in one, namely, the Committee on Information and Communication (CIC).

- *The Committee on Applications and Technology of the Information Infrastructure Task Force (IITF)*. The IITF was created by the president to propose the policies and initiatives needed to accelerate the deployment of the National Information Infrastructure and assist in the implementation of those policies. The OSC participates in the IITF's Applications and Technology committee as well as its Government Information Task Services Working Group, its Technology Policy Working Group, and its Government Information and Technology Services (GITS) subcommittee. The OSC is also a member of the GITS Electronic Mail Subgroup (GEMS), whose purpose is to ensure interoperability between Government agencies and the research, development, and educational programs that use e-mail directory services.
- *The High Performance Computing and Communications Information Technology (HPCCIT) subcommittee of the NSTC Committee on Information and Communication*. This subcommittee addresses issues that relate to the retention of U.S. leadership in the application of computing communications technologies. The HPCCIT subcommittee is particularly concerned with government policies that affect advanced segments of industry and national security agencies. It is composed of several working groups, including the Science and Engineering Computing Working Group, the Education Working Group, and the Network Working Group.
- *The Federal Network Council (FNC)*. The FNC is chartered to establish a long-term strategy for the operation and evolution of the Internet and other national computer networks. The Council is particularly concerned with the support of research and education.
- *The National Academy of Sciences (NAS)*. The OSC participates in NAS's board on Mathematical Science and its Computer Science Technology Board.
- *The Coordinating Committee for Informatics Research, Development, and Application (CCIRDA)*. CCIRDA functions as the chief planning and coordinating body for informational, computational, and communications requirements throughout ER.
- *The Federal Interconnectivity Requirements Panel (FIRP)*. The OSC played a key role in establishing and funding this panel. One of its first actions was to review Standard 146 of the Federal

Information Processing Standards (FIPS), which contained provisions that would have prohibited the use of TCP/IP protocols by federal agencies. As a result of the FIRP's review, the networking provisions of the FIPS were modified to permit the use of TCP/IP protocols.

Liaisons Maintained by ESnet Management

ESnet management generally provides the technical and administrative liaisons necessary for ESnet to participate in national and international Internet activities. These liaisons are essential to maintain many important elements of an internationally distributed network. Some of the current liaison bodies are:

- *The Internet Engineering Task Force (IETF)*. The IETF is an open forum whose purpose is to engineer solutions to problems that affect network operations in the short-to-medium term. Working groups are formed to deal with specific aspects of such problems.
- *The Federal Engineering and Planning Group (FEPG)*. The FEPG provides a forum within which federal network providers can share information and address networking issues of common concern.
- *The North American Network Operations Group (NANOG)*. This group meets 3–4 times yearly and provides a forum in which managers of U.S. Internet networks can interact and respond to issues of common concern.

Liaisons Maintained by National Laboratories and Universities

The national laboratories and many universities maintain numerous liaisons with industrial contacts and independent links to regional networks. Many of these links intersect with the interests of ESnet in ways that benefit not only ER but the wider DOE community and the nation as a whole. Included among these liaisons are:

- *Industrial Contacts*. Historically, the national laboratories and universities have used networking as a means to facilitate the geographical dispersion of collaborative science. The sciences have become dependent on high-performance networking, which many consider an essential medium for the conduct of modern science. This dependence has led laboratories and universities to work directly with industry in developing and testing new network products and capabilities.

- *Regional Networking Contacts.* Several ESnet-related laboratories and universities connect directly to regional and state networks. These connections provide ESnet with independent pathways to all of the systems within the regional networks in question. These independent pathways will play an important role in the transition from an NSF-provided Internet to commercial service.
- *Standards Development.* Laboratory and university personnel participate in the standards committees that provide the foundations for the development of new communications technologies.

ESnet's International Activities

ESnet and its participating sites have assisted in establishing initial international connectivity for a number of countries, including Italy, Germany, Japan, Brazil, and, most recently, China. ESnet and its constituents actively participate in several international organizations and activities including:

- *The Coordinating Committee for International Research Networking (CCIRN).* The mission of CCIRN is to stimulate cooperative international research by promoting enhanced, interoperable networking services. CCIRN also seeks to: (1) coordinate international connections between the networks represented on the committee, (2) promote the evolution of an open international research network, (3) coordinate the development of international network management techniques, and (4) disseminate the results of networking research and development.
- *The International Engineering and Planning Group (IEPG).* ESnet also participates in the forums provided by the IEPG.

ESnet Today

The Energy Sciences Network is a nationwide data communications network managed and funded by the U.S. Department of Energy Office of Energy Research (DOE/ER) for the purpose of supporting multiple-program, open scientific research. ESnet fulfills this purpose by providing widespread access to research facilities, state-of-the-art communications between collaborators, and on-line information services and other related services.

How is ESnet Managed?

ESnet is an ongoing activity whose success depends heavily on the cooperation and collaboration of the ER community. Although responsibility for the implementation and operation of ESnet resides with ESnet management at NERSC, guidance for the network comes from two key sources, the ESnet Steering Committee and the ESnet Site Coordinating Committee. The Steering Committee defines program requirements and recommends directions and priorities to the implementation staff, while the Site Coordinating Committee represents

the institutions that benefit from the use of ESnet and assumes responsibility for disseminating information to them.

Who May Use ESnet?

Currently, ESnet's principal mandate is to support the programs sponsored by DOE/ER. The major OER-supported programs are: Basic Energy Sciences, High Energy Physics, Magnetic Fusion Energy, Nuclear Physics, Health and Environmental Research, and Applied Mathematical Sciences. ESnet usage in support of other activities, such as interagency collaboration or foreign-country access, may also be authorized.

Since its inception, ESnet has been dedicated almost exclusively to the Energy Research community. However, recent years have seen a rapid growth in the use of networked data communications services by researchers, scientists, and other personnel throughout DOE. With the increasing integration of network services into the workstyles of a broad spectrum of personnel, people from a number of other DOE program areas have expressed interest in availing themselves of ESnet's services. Such usage must first be formalized through an internal agreement known as an intra-agency Memorandum Of Understanding (MOU) between the group in question and the Office of Scientific Computing. MOUs are currently in place between the OSC and IRM, EM, and HR. An MOU between OSC and DP (Defense Programs) is currently under consideration.

The Network Today

The ESnet Backbone

ESnet's initial T1 (1.3–1.5 Mbps) backbone became fully operational in early 1990. Today there are more than 30 sites directly connected to the backbone, and portions of it have been upgraded to T3 links, which provide bandwidths of 45 Mbps (see Figure 6-1). Connectivity to numerous other scientific and educational locations is provided through extensive interconnections with other networks that comprise the global Internet. In addition, ESnet supports multiple network-level protocols, including the Department of Defense Internet Protocol (DOD-IP), DECnet Phase IV, and the Open Systems Interconnection ConnectionLess Network Protocol (OSI CLNP).

Figure 6-1. ESnet backbone as of mid-1994

Network Capabilities

Networking data communications can be characterized in terms of two important parameters, performance and connectivity. Performance metrics include both bandwidth and latency, while connectivity defines the geographical extent of the destinations that can be reached over the network. ESnet's current project implementation seeks to address both network performance and connectivity, in response to program requirements ultimately based on the capabilities that end users require for the support of their research.

Performance

When ESnet's T1 topology first became operational in early 1990, it was a trunking system based on point-to-point 1.5 Mbps circuits with less than three "hops" (i.e., interconnecting links) between sites, on average. In early 1990 this configuration provided what was considered to be a respectably high level of performance, particularly for a national wide-area network (WAN). However, by early 1991 it was becoming clear that an upgrade of capability would be required within a few years, and planning for a T3 (45-Mbps) capability was started. An ESnet review in early 1992 recommended that additional attention be given to this effort.

Simultaneously, a national program to enhance American competitiveness got underway. This program, the High Performance Computing and Communications (HPCC) Initiative, incorporated a network focus on similar capability, with a target of gigabit-per-second networking by sometime in 1996.

ESnet was thus confronted simultaneously by two requirements with rather differing implications: the requirement to provide production-quality network services, as specified by the existing OSC programs, and the requirement to promote leading-edge technology mandated by the HPCC. In response to these divergent requirements, a plan was developed to establish a cooperative approach with a vendor whereby ESnet would make use of emerging communications technology in a carefully implemented manner while serving a large and sophisticated customer base to help shake down the new technology. A Request For Proposal released in February 1992 described the details of this plan.

In June of 1992, a vendor was selected to supply communications service based on T3 ATM (Asynchronous Transfer Mode) technology. By mid-1994, T3 links were in place spanning the nation from PPPL (Princeton Plasma Physics Laboratory) to NERSC at LLNL, as shown in Figure 6-1. At that time, many of the T1 lines had become very heavily loaded, due to the dramatic growth in ESnet's traffic (see Figure 6-2), and the conversion of the remainder of the backbone began.

Figure 6-2. The growth of ESnet's traffic, 1990-94

Connectivity

Connectivity requirements come in many forms: connections of new sites, general requirements to gain access to University and other collaborator sites, requirements to connect with regional networks, requirements to share access with other agencies, and requirements to provide international access to selected sites or countries. This interconnectivity is provided in a variety of means that include indirect connections through other R&E networks (see Figure 6-3) in addition to direct connection of sites to the ESnet backbone.

Figure 6-3. ESnet's connectivity to regional networks

As of late 1993, the sites with direct connections to the ESnet backbone were:

ANL:	Argonne National Laboratory, Argonne, IL
BNL:	Brookhaven National Laboratory, Upton, NY
CIT:	California Institute of Technology, Pasadena, CA
CEBAF:	Continuous Electron Beam Accelerator Facility, Newport News, VA
DOE:	DOE Office of Energy Research, Germantown, MD
FIX-E:	Federal Interagency eXchange-East, University of Maryland, College Park, MD
FIX-W:	Federal Interagency eXchange-West, NASA Ames Research Center, Mountain View, CA
FNAL:	Fermi National Accelerator Laboratory, Chicago, IL
FSU:	Florida State University, Tallahassee, FL

GA:	General Atomics, La Jolla, CA
GRM:	Grumann Aerospace, Princeton, NJ
HARVARD:	Harvard University, Cambridge, MA
ISU:	Iowa State University, Ames, IA
ITER:	International Thermonuclear Experimental Reactor Project, La Jolla, CA
LBL:	Lawrence Berkeley Laboratory, Berkeley, CA
LLNL:	Lawrence Livermore National Laboratory, Livermore, CA
LANL:	Los Alamos National Laboratory, Los Alamos, NM
MIT:	Massachusetts Institute of Technology, Boston, MA
NEVIS:	Nevis Laboratories, Columbia University, Irvington, NJ
NYU:	New York University, Upton, NY
OER:	US DOE, Office of Energy Research, Germantown, MD
ORAU:	Oak Ridge Associated Universities, Oak Ridge, TN
OSTI:	Office of Science and Technology Information, Oak Ridge, TN
ORNL:	Oak Ridge National Laboratory, Oak Ridge, TN
PNL:	Pacific Northwest Laboratory, Richland, WA
PPPL:	Princeton Plasma Physics Lab, Princeton, NJ
SNLA:	Sandia National Laboratory, Albuquerque, NM
SNLL:	Sandia National Laboratory, Livermore, CA
SLAC:	Stanford Linear Accelerator Center, Palo Alto, CA
SSC:	Superconducting Super Collider, Waxahachie, TX
UCLA:	University of California, Los Angeles, CA
UTA:	University of Texas, Austin, TX

The broad regional connectivity illustrated above has been developed in part as a strategic response to NSF's announced plans to withdraw from providing general networking support for the US academic and research community. ESnet has established, or plans to establish, direct "peering" relationships with nearly all US regional networks.

The 18 months between mid-1993 and early 1994 saw the emergence of a major set of requirements for increased connectivity to Europe, particularly to Germany and Italy, and to the former Soviet Union, particularly to Russia. An agreement was established with DFN (German Research Network) in Germany and INFN (Italian National Institute for Nuclear Physics Network) in Italy to jointly fund a T1 circuit to Germany and a 512 kbps circuit to Italy. Orders for these "half-circuits" were placed and installation was completed in mid-1994.

Establishing connectivity to Russia is proving much more difficult. An effort to collect requirements has produced a list of approximately 35 sites of interest in Russia, and new sites are being requested at the rate of several per month. A trip was made in August-September 1993 to St. Petersburg and Moscow to better assess possible approaches and to develop contacts and working relationships with Russian counterparts. Development of plans to provide connectivity to both of these cities has

been identified as a first step toward general access into Russia and other former Soviet Union countries.

On a shorter time-frame, ESnet has begun providing indirect access to available Russian institutions through other networks, such as NSFnet. These connections are typically of fairly low performance but do offer at least an initial opportunity for network connectivity. In December 1993, we established a connection to the prototype GIX (Global Internet eXchange) located in the Washington, D.C., area to provide some additional (albeit indirect) connectivity.

ESnet Services

Many people think of ESnet primarily in terms of its ability to deliver data packets. Although the mainstay of ESnet services certainly is its ability to provide data communications on a national and international basis, the entire set of ESnet services is much more extensive. These services can be classified as follows:

- Network Operations Management
- Network Infrastructure Services
- Information Services
- Videoconferencing Services.

These services are described in detail in the following subsections.

Network Operations Management

The ESnet staff currently manages over 45 network routers deployed on a national basis. Protocols supported include Department of Defense Internet Protocol, DECnet Phase IV, and the Open Systems Interconnection Connectionless Network-layer Protocol. Bandwidth supported is typically 1.5 Mbps, although several 45 Mbps links are in operation. Broad interconnectivity is supported to other agencies, regional and commercial networks, and international entities.

There are a number of unique challenges involved in providing the communications infrastructure support for a nationwide backbone network. A data communications network is inherently a distributed system spanning a geographic area that may range from a room or an office to the entire globe. The Internet in particular poses a difficult environment in which to operate a networking service capable of providing production-quality support to ER research activities. (The world-wide DECnet poses a number of similar problems, and a few unique ones of its own.) The Internet is essentially an interconnected, distributed system of autonomous networks that are also distributed systems. Each of these autonomous networks is termed an *autonomous*

system (AS), as it interconnects with other ASs in both an ad hoc and a sovereign manner. Production operation of the Internet therefore depends on the successful interoperation of an interconnected collection of such autonomous systems. A failure or an independent action in one AS can easily have an unanticipated impact on others.

In this environment, maintaining production-quality service requires full-time monitoring and problem resolution, and ESnet is operated 24 hours a day, 7 days a week. The ESnet Network Operations Center (NOC), shown in Figure 6-4, provides 24 hour-per-day monitoring and control capabilities for the various network components that comprise ESnet. The Operations Center is operated by the ESnet staff and NERSC supercomputer operations personnel. The NOC staff continuously monitors ESnet's backbone facilities to verify the network's integrity and to gather routine statistics that will facilitate troubleshooting and long-term planning. Electronic mailboxes for network information, network operations, and trouble calls are provided. An on-line trouble ticket system exists to ensure that all reported problems will be properly tracked. Requests for information about ESnet can be directed to info@es.net. Problems can be reported to trouble@es.net.

Figure 6-4. The ESnet Network Operations Center

Maintaining high network availability depends not only on problem resolution but also on constant planning, proactive implementation, and the management of critical transitions in carefully planned stages. For example, a number of schemes that attempt to define a next-generation Internet Protocol (IPng) were being defined and evaluated as this Program Plan went to press. However, the time scale required for such a major transition made a shorter-term solution essential. An interim proposal called Classless Inter-Domain Routing (CIDR) aimed to ease the demand on router resources by aggregating networks into address blocks. As of this writing, ESnet had just completed a major transition to an inter-AS routing capability called BGP-4 (Border Gateway Protocol, version 4), which is capable of supporting the aggregated addressing associated with CIDR by "tagging" network addresses. The implementation of these protocols was a major change that could not have been done piecemeal. Following several months of planning, a group of ESnet people totaling as many as eight at any one time completed the transition to BGP-4 in a single all-night session, with only minimal impact on users.

Network Infrastructure Services

Internet Protocol Domain Name Service (IP/DNS)

ESnet provides Internet Protocol Domain Name Service to the ES community through two separate name servers, DNS-WEST and DNS-EAST. These name servers provide the primary IP address translations for the es.net domain. In addition, either server can be queried to resolve IP address translations for any IP host on the Internet.

In addition, these machines provide secondary name service for many of the ESnet sites. Employing this functionality, sites can off load many name server queries destined for their local name servers to DNS-WEST and DNS-EAST. This procedure has the advantages of providing a backup to a site's name server while reducing its load.

DEC Distributed Naming Service (DECdns)

DECdns is a DECnet/OSI network-wide service that makes it possible to use network resources without knowing their physical locations. Users and applications can assign DECdns names to resources such as nodes, disks, and files. The creator of the name also supplies other relevant information (such as the resource's network address) for DECdns to store. Users then need only supply the name, and DECdns will act as a lookup service, providing the rest of the data when necessary.

The OMNI namespace has been created to support the transition to DECnet Phase V within the DECnet Internet. ESnet supports two root DECDns servers within the ESnet portion of the DECnet Internet (ESnet-DECnet). As ESnet sites bring up DECDns servers, they are added to the OMNI namespace.

Distributed Time Services

The purpose of the Internet's Network Time Protocol (NTP) is to synchronize the clocks on multiple networked hosts. Redundancy and sanity checking are features of the protocol. ESnet provides several stratum-1 time servers, each with a radio clock served by the U.S. time standard broadcast from Fort Collins, CO, by NIST (National Institute of Standards and Technology). NTP software is available from ESnet in the directory Public-Domain.

Authentication Services (Under Development)

ESnet operates an experimental Kerberos Version 5 Authentication server. Kerberos provides a strong form of user identification, and use of Kerberos can eliminate the need to send passwords over the Internet.

Information Services

The Network Information Center (NIC)

ESnet provides a number of information and support services, which are described below. These services are provided via a pair of duplicate servers in ESnet's Network Information Center, which is located at the Network Operations Center at NERSC. In general, users must have access to appropriate client software on a local PC, MAC, or workstation to use these services, although many services are available via telnet or ftp. In some cases, remote-access client capability is also provided on the NIC, with the user accessing the client software via telnet or X-windows. In still other cases, remote client software is made available on the ESnet NIC for users to download and use on their local computing platforms.

Users can gain direct access to the ESnet NIC server through a variety of network applications, from ftp and telnet to WAIS (Wide Area Information Server), Gopher, and Mosaic (see Table 6-1).

The ESnet Network Information Center servers are the primary information repository for the ESnet community. The information available on these servers includes ESnet committee documents, such as meeting notes and notices for the ESSC and the ESCC and associated subcommittees, in addition to ESnet mailing lists, NetNews archives, public-domain software, and files and documents regarding the latest

networking technology and research topics All information on the ESnet NIC servers is full-text WAIS indexed.

Table 6-1. Methods of Accessing ESnet's Network Information Center Server

APPLICATION	ADDRESS
Mosaic	http://www.es.net
Gopher	gopher://gopher.es.net
ftp	ftp://ftp.es.net
DECnet	esnic:: (42158::)
NFS	nfs.es.net
AFS	/afs/es.net/nic
e-mail	nic@es.net

Directory Services

OSI X.500 Directory Service (White Pages)

The White Pages Pilot Project is part of a world-wide Distributed Directory Service pilot project based on the X.500 standard. More than 30 countries participate in the global pilot project, which contains more than 1 million entries. The pilot Directory Service contains information primarily about persons and organizations. The information in the Directory is hierarchically structured, with entries arranged in the form of a tree. Users can browse this tree or retrieve information about a person, generally by knowing that person's name and organization.

The ESnet community is well-represented in the White Pages pilot project. Currently, there are more than 15 backbone sites represented in the Directory, which comprises over 100,000 entries. These entries, which are primarily descriptions of persons, include e-mail addresses, postal addresses, telephone numbers, and other optional descriptive information. This information can be accessed in a variety of ways, including the use of X.500 clients, finger, whois, e-mail, Gopher, or WWW.

Ph (Phone Book)

Ph (Phone book) is an fully indexed, fast-access on-line white-pages directory service developed and freely distributed by the Computer and Communications Services Office at the University of Illinois at Urbana. Ph primarily contains information about people, such as e-mail addresses, postal addresses, telephone numbers, and other descriptive information. There are over 180 Ph Servers on the Internet today.

Electronic Mail Services

OSI X.400 and SMTP

ESnet's Message Handling Services (electronic mail) provides fully compliant SMTP (Simple Mail Transfer Protocol) and ISO/OSI X.400 mail transport services. A gateway between SMTP mail and X.400 mail is provided to support full, bidirectional connectivity between these two communities. ESnet's mail system supports mail list archiving and NetNews cross-posting, as well as remote list management. All archives are full-text WAIS indexed and are accessible via ESnet's WWW, Gopher, and WAIS services. Remote list management is accomplished via an Xwindow tool that allows "list owners" to build and manage their own ESnet mailing lists.

NetNews and Bulletin Board Services

NetNews is a global service that permits a group of people with a common interest to easily exchange articles tagged with one or more universally recognized labels, which are called NewsGroups. There are public NewsGroups for almost every area of interest, from hobbies to scientific research.

ESnet's NetNews service carries many of the public NewsGroups available on the Internet as well as NewsGroups local to the ESnet community. ESnet's NetNews service supports e-mail cross-posting for any mailing list on the ESnet mail system. The NetNews service provides for NewsGroup archiving on the ESnet Information Server. The NewsGroup archives are full-text WAIS indexed and are accessible via ESnet's WWW, Gopher, and WAIS Services.

Videoconferencing Services

Recent advances in technology have made the cost of videoconferencing low enough—in terms of both dollars and communications bandwidth—that it is growing rapidly in usage and deployment. Two basic modes of videoconferencing are currently in use within ER, ISDN-based conference-room videoconferencing (see Figure 6-5) and packet-based desktop videoconferencing.

Conference-Room Videoconferencing

The Video Conferencing Service (VCS) provides a full range of ISDN-based videoconferencing services to the ESnet user community. Users can check for available time slots and book conferences by accessing the Video Conferencing Service Scheduler (VCSS), an on-line automated reservation system. The VCSS tracks available resources,

Figure 6-5. Conference-room videoconferencing

schedules conferences, generates reports, and sends electronic mail notification. The Video Conferencing Service also provides the bridging equipment, multiplexers, and ISDN circuits necessary to host multiway videoconferences. In addition, a Help Desk has been established to assist VCS users with operational issues. Other video conferencing services

provided by this group include support of open standards and an interoperability gateway for conferencing with outside organizations.

The ISDN-based architecture currently in use within the ER community employs VTEL videoconferencing equipment in ordinary meeting rooms. Communications requirements are met via point-to-point circuits, either dedicated or dial-up, operating at speeds of 128 kbps and up. In order for two sites to share a two-way videoconference, a communications channel must be established between them for the duration of the conference. Multiway videoconferences (i.e., with three or more participants) require the use of an MCU (Multiway Conference Unit) that essentially acts as a meeting chairman, deciding who is "given the floor," that is, who is allowed to speak to and be seen by the other parties at any given time during the conference.

Workstation-Based Videoconferencing

Desktop videoconferencing uses the participant's workstation monitor for display of the incoming video, while the outgoing video is picked up by a small camera that is typically mounted on or near the monitor. The video and audio data are "packetized" into IP packets and "multicast" over the Internet/ESnet. Within the ER community, the most commonly used desktop video facilities employ a combination of publicly available software packages developed in part at LBL.

A multicast is basically a selective broadcast in which data are sent from one source to specified multiple destinations in a way that minimizes redundant traffic over the intervening links. To achieve this end, a virtual network called the *Mbone* (for *Multicast backbone*) has been constructed on top of the existing IP network. ESnet implemented its own segment of the Mbone during 1993, constructing multicast routers from workstations and locating them at strategic locations around the ESnet perimeter.

The ESnet Mbone is currently an experimental service provided to the community to make it possible to share multicast data. Present uses of this technology allow audio, video, images, and raw data to be sent simultaneously to many users worldwide. However, as of this writing wide-area desktop videoconferencing still has some significant performance problems. Frame rates as low as a few seconds per frame and frequent audio breakups are still typical. Another problem is that packet-based videoconferencing generates an enormous amount of traffic on a steady basis. For ESnet, this creates a very large traffic load on an already heavily loaded network. Even with the sophisticated compression schemes being used, a single multicast "meeting" may generate more than 100 kbps of traffic continuously for hours.

While desktop videoconferencing is still an experimental service, its deployment is growing steadily, and it promises to become a mature offering within the next several years. ESnet-wide desktop videoconferences have been held, allowing participants to contribute to meetings without ever leaving their offices. In addition, the ESnet Mbone is connected to the larger Internet Mbone, giving our users the capability to interact in real time with colleagues around the globe.

High Energy Physics

*Larry Price
High Energy Physics Division
Argonne National Laboratory*

High Energy Physics is concerned with the structure of matter and forces at the most fundamental level. Paradoxically, the quest to understand ever-smaller and more basic components of matter has required particles of ever-higher energy. Apart from the theoretical component, most HEP investigations are carried out at the major accelerator centers. In the U.S., these centers are the Fermi National Accelerator Laboratory, the Stanford Linear Accelerator Center, Brookhaven National Laboratory, and Cornell University's Wilson Synchrotron. 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 Institute for High Energy Physics (IHEP) in Protvino, Russia. With the termination of the SSC Laboratory, U.S. HEP researchers concerned with the search for the Higgs boson will look toward a major effort at the Large Hadron Collider (LHC) now under construction at CERN. At the same time, the *B* Factory under construction at SLAC (see sidebar) will be a significant new U.S. HEP site, and a community of physicists interested in studying matter-antimatter asymmetry and charge-parity violation has begun work on a detector for that facility.

The *B* Factory at SLAC

On January 11, 1994, Secretary of Energy Hazel O'Leary attended the inauguration of the *B* Factory construction project at the Stanford Linear Accelerator Center. The new facility is designed to produce copious quantities of *B* mesons and anti-*B* mesons, with the aim of enabling physicists to test certain aspects of the Standard Model of particle physics. The ultimate goal is to gain insight into one of the central mysteries of cosmology, namely, that of why we observe such a preponderance of matter over antimatter in the universe today. Current theory holds that matter and antimatter existed in equal quantities at the moment of the Big Bang. Given that initial condition, we should not observe a universe composed largely of matter. Instead, one of two conditions should obtain today: Either matter and antimatter should have annihilated each other, or antimatter should exist in quantities roughly equal to those of matter. The current explanation for the matter-dominated universe is that there was some small preference for matter over antimatter in the original process of creation. If we can assume that physical processes have operated consistently since the Big Bang, then the key to understanding the preference for matter creation must lie in the phenomenon of charge-parity (CP) violation. Given the current state of their art, physicists believe that CP violation is best studied by examining the decay of *B* mesons and their antiparticles, an investigation that requires the production of billions of *B* mesons and anti-*B* mesons. Hence the need for a *B* Factory.

The new *B* Factory will actually be an upgrade of SLAC's existing PEP (positron-electron project) facility. The construction project will entail renovating the existing high-energy PEP storage ring, adding a new low-energy storage ring inside the same PEP tunnel, and creating and installing a detector designed expressly for CP-violation studies. Data collection is scheduled to commence in the spring of 1999. After that point, the project is expected to generate approximately 100 terabytes of raw, reconstructed, and Monte Carlo data each year. Large amounts of this data must be distributed within a collaboration that includes nearly 400 physicists representing 75 institutions in nine countries (the U.S., Canada, the UK, France, Italy, Germany, Russia, the PRC, and Taiwan). Wide-area networking will play a crucial role in this widely-distributed enterprise. Essential interactive services will include e-mail, NetNews and World Wide Web access, videoconferencing (workstation-based as well as conference-room-based), and remote log-in using X windows. Remote operation of parts of the detector system will be made available to some external sites for problem diagnosis and debugging. In addition, distributed software development will require a system for coordinating code development and automating its distribution. Finally, the prompt data analysis needed to support timely publication of experimental results will require the transfer of large data sets between institutions. Data transfers on this scale could easily overwhelm today's network capacities. However, they also offer the promise of a data transport medium that is faster, more reliable, and less labor-intensive than traditional tape distribution.

Above, B Factory layout; inset, Hazel O'Leary at B Factory groundbreaking in January of 1994

Computing and Networking in HEP

Experiments at the major accelerator centers are large-scale enterprises, typically involving 100 to 500 physicists and at least as many engineers and technicians during the construction phase. From initial conception to final data collection, such experiments range in duration from five to fifteen years. In the next generation of experiments, the two large collaborations focused on the LHC will include about 1500 members each. These collaborations will involve participation by multiple institutions, with an international mix.

The primary factor determining participation in a collaborative HEP experiment is an interest in a certain approach to physics; there is little regard for geographic proximity in the formation of such a collaboration. Good communications—particularly via computer networking and videoconferencing—are therefore critical in enabling a collaboration to function at all, and communications must be extremely good if the collaboration is to function smoothly. Computer networking is especially crucial in HEP experiments because their complexity requires the use of computers at every stage of operation. For example, large codes are written to acquire, store, and analyze large samples of data, and each of these processes will typically involve collaborators at widely separated institutions. Fast, reliable, sophisticated networks are indispensable to support such joint efforts.

Future Requirements

Network Services

The basic services available via computer networking have been integrated into the operation of current HEP experiments and into the planning of the next generation of experiments. These services also come into play in some phases of theoretical work. To serve these purposes, such services as electronic mail, file transfer services, virtual terminal service, remote access to files, and remote job submission and job monitoring must operate with complete reliability and at high speeds. The entire HEP research community requires access to these services via TCP/IP, and a significant minority of HEP physicists continues to require access to them via DECnet. The long lifetime of HEP experiments means that DECnet support will be required for this subset of users through the year 2000 and perhaps beyond.

In 1990, HEP groups began an experiment in the use of videoconferencing for scientific collaborations. With an initial link established between LBL and SLAC (soon extended to the SSC Laboratory),

conference-room videoconferencing was found to be a highly effective medium for collaborative meetings. By 1993, HEP usage of videoconferencing was extensive enough that time slots were difficult to find in the normal workweek, including the early-evening slot most convenient for conferences involving Japan. By then, HEP's videoconferencing system had grown to 17 sites, two of which were in other countries. In 1994, the HEP system was integrated into ESnet's Video Conferencing Service, and a transition was made from multiplexed use of ESnet's leased lines to use of circuit-switched video over the commercial networks. As videoconferencing evolves technically, its use will continue to be of great advantage to HEP and will expand as fast as network bandwidth permits. It seems likely that within two years, as much bandwidth will be devoted to video as to data.

At the time of this writing, there was a great deal of pent-up demand for conference-room video service. The extent of this demand would suggest that the this service will become much more widely used as its cost continues to fall. However, the lack of universal operability poses an additional barrier to increased usage of this service. This problem stems from a lack of interest on the part of the telephone companies in ensuring that the services introduced after the breakup of the nationwide Bell system are as interoperable as those introduced before that time. The lack of universal interoperability is currently a problem for switched data services (including conference-room videoconferencing), and it seems likely that it will be an equal problem for the new services envisaged as parts of the "information superhighway." ESnet's VCS is currently supplying interconnection services in cases where telephone companies are unwilling or unable to ensure interoperability. Considering the projected visions of the "information superhighway" and the failure of the telephone companies to address the interoperability problem, it seems likely that ESnet will need to extend such interconnection services.

Desktop, or workstation-based, videoconferencing is an alternative service currently under development by a number of companies. As its functionality improves and as standards are defined for its implementation, HEP collaborations can be expected to make heavy use of workstation videoconferencing. However, the usage model incorporated into current development plans for this style of videoconferencing may pose a significant barrier to its widespread use. This model assumes that workstation-based videoconferences will involve only a few participants, each using a separate workstation. The resulting implementations of workstation video will be suitable for conferences between a few widely-separated individuals but will not support meetings involving up to 20 people at any one site. For videoconferences that do involve numerous participants at a single site, it is strongly recommended that the larger group(s) meet in a conference room, where they can interact more fully

among themselves, even though they are communicating with others at remote sites via workstation-based video.

Other new services are needed urgently. One such service is a complete, readily accessible directory of institutions and individuals. At present, the lack of such a directory service significantly reduces the utility of the network. ESnet management and the ESCC have made a significant start toward solving this problem, but the emerging directory tools are not widely known within the Energy Sciences community. The World Wide Web system, which was developed by the European HEP community, promises to facilitate the use of directory services and many other kinds of information services over the ESnet. To cite another emerging requirement, HEP's use of graphical windowing (via X-11, Motif, etc.) to access remote computers is becoming routine and can be expected to grow rapidly. Because of this demand alone, bandwidth requirements for interactive use will grow by an order of magnitude over the next 2-3 years.

Connectivity

HEP networking needs are largely defined by the requirements of large accelerator-based experiments. However, the smaller component of HEP research that is not accelerator-based creates its own distinctive subset of requirements. These experiments typically seek to detect cosmic rays or radiation from rare, spontaneous terrestrial events. Such experiments involve shielding the detector system within a mountain or deep in a mine to ensure the sensitivity required to isolate a tiny signal. This aspect of HEP work requires efficient, reliable network communications to such remote places as Soudan, Minnesota, Dugway, Utah, the Gran Sasso Tunnel in Italy, or the Baksan mine in Russia.

HEP groups at laboratories and universities require reliable, universal connectivity among themselves. At the same time, the importance of access to major experiments and databases at Fermilab, SLAC, BNL, CERN, DESY and China's IHEP creates a requirement for higher-bandwidth connections to those sites. The table at the end of this section lists the sites of major HEP experiments that U.S. institutions currently participate in; foreign participation is also listed for each collaboration. As noted above, HEP requires fully capable network connections not only between each of the participating institutions and each of the experimental sites but also between all of the participating institutions.

Because the HEP field and its collaborations are highly international in scope, connections to major locations in Europe and Japan need to be as good as domestic connections. The fact that this is not true today clearly limits the efficiency and productivity of present HEP collaborations. The major requirements for international links during the remainder of the 1990s are as follows:

- A 1 Mbps link to CERN in Geneva, Switzerland, is needed now. As the LHC program becomes established, this link should grow steadily in capacity to 10 Mbps by the end of the decade.
- A 0.5-Mbps link is needed to DESY in Hamburg, Germany. This link should grow in capacity to 1.5 Mbps by the end of the decade.
- A 1-Mbps link is currently needed for general connectivity to the rest of Europe. This link's capacity should grow to 10 Mbps by the end of the decade.
- A 1-Mbps link is now needed to KEK in Tsukuba, Japan (and to the rest of Japan); this link's capacity should grow to 10 Mbps by the end of the decade.
- A 0.5-Mbps link is now needed now to IHEP, in Beijing, China, with connectivity to other HEP institutes in China. This link's capacity should grow to 1.5 Mbps by the end of the decade.
- A 0.5-Mbps link is now needed to key institutes in Mexico, Brazil, and Argentina, with connectivity to other parts of Central and South America. This link's capacity should grow to 5 Mbps by the end of the decade.
- A 1-Mbps connection is now needed to JINR (Dubna), IHEP (Protvino), and other points in the Moscow Region in Russia. Connectivity is also needed to institutes in St. Petersburg and Novosibirsk and to institutes and universities throughout Russia and other former Soviet republics.

Performance

Estimated bandwidth requirements for new international connections have been indicated above. A highly accurate method for estimating HEP's domestic bandwidth requirements was employed by the HEPnet Review Committee (HRC) in its 1988 report on HEP computer networking.¹ That analysis led to the conclusion that ESnet's then-current 56 kbps X.25 backbone lines would become saturated by early 1989 and to the recommendation that plans should be made for an almost immediate upgrade of the backbone bandwidth to T1 speeds (1.5 Mbps).² Section III of the HRC report further concluded that planning for the next step in bandwidth above T1 speeds should begin sometime in 1990.

The subsequent upgrade of the ESnet backbone to T1 lines was completed in time to keep pace with HEP's bandwidth requirements, validating the HRC report's estimate. However, the step up to the next bandwidth level (T3, or 45 Mbps, speeds) has not occurred fast enough to meet the accelerating HEP demand. As of this writing, many lines used by HEP institutions are now becoming saturated, and some emerging uses of the ESnet are on hold until the upgrade to T3 lines can be completed. Full implementation of the T3 backbone is therefore a critical priority.

Planning and Management

As we have stressed, adequate networking support is a critical requirement in all phases of HEP research, both experimental and theoretical. Prompt network access, adequate bandwidth, and essential network services are fundamental requirements for all HEP researchers. In addition, ESnet must provide sufficient network management resources to prevent interruptions of service. ESnet's management must also be able to forecast requirements well enough to provide the necessary performance and connectivity before their lack hampers the scientific program. Finally, it is crucial for management to keep in mind that the growth in demand for network services is fueled by the emergence of qualitatively new capabilities as well as by quantitative increases in usage of existing capabilities. Both trends must be tracked and taken into account in the planning of further ESnet development.

Notes

1. L. E. Price, et al., *High Energy Physics Computer Networking: Report of the HEPnet Review Committee*. DOE/ER-0372 (1988), pp. 38-42.
2. *ibid.* p. 43.

Major HEP Experiments and Collaborating Institutions

Beijing Spectrometer

Beijing, Inst for High Energy Phys
 Boston U
 Cal Tech
 Colorado State U
 Hawaii U
 Mit, LNS
 SLAC
 Texas U, Dallas
 UC, Irvine
 Washington U, Seattle

BNL-865

Basel U
 Brookhaven
 Dubna, JINR (Russia)
 Moscow, INR (Russia)
 New Mexico U
 PSI, Villigen, Switzerland
 Pittsburgh U
 Tbilisi State U (Georgia)
 Yale U
 Zurich U

BNL-871

UC, Irvine
 Stanford U
 Texas U
 William And Mary Coll
 Richmond U

BNL/RHIC Phenix

Brasil U Sao Paolo
 McGill U
 CIAE
 IHEP
 Inst Mod Phys
 Peking U
 U Munster
 BARC Bombay
 Hiroshima U
 INS U Tokyo
 KEK
 Kyoto U
 Nagasaki Inst
 Nat Inst Rad Sci
 U Tokyo
 U Tsukuba
 Chung ang U
 Korea U
 Seoul Nat U
 Soong Sil U
 IHEP Protvino
 INR Moscow
 ITEP Moscow
 JINR Dubna
 Kurchatov Inst
 PNPI St Petersburg
 Lund U
 BNL
 Columbia U

BNL/RHIC Phenix, cont'd.

Florida State
Georgia State
Idaho NEL
Iowa State Ames Lab
LANL
LLNL
Louisiana State
MIT
ORNL
SUNY Stony Brook
UC Riverside
U Alabama
U Tennessee
Vanderbilt U
Yale U

CERN Aleph

Annecy, LAPP
Barcelona, Autonomia U
INFN, Bari And Bari U
Beijing, Inst High Energy Phys
CERN
Clermont-Ferrand U
Bohr Inst
Democritos Nuclear Research
Center
Ecole Polytechnique, LPNHE
Edinburgh U
Florence U And INFN, Florence
Florida State U, Scri
Frascati (Italy)
Glasgow U
Heidelberg U, IHEP
Imperial Coll, London
Innsbruck U
Lancaster U
Mainz U, Inst Phys
Marseille, Cppm
Munich, Max Planck Inst
Orsay, LAL
Pisa U And INFN, Pisa And
Pisa, Scuola Normale
Superiore
Royal Holloway - Bedford Coll
Rutherford
Dapnia, Saclay
UC, Santa Cruz
Sheffield U
Siegen U
Trieste U, Ist Fis And INFN, Trieste
Wisconsin U, Madison

CERN Delphi

NIKHEF, Amsterdam
Antwerp U
Democritos Nuclear Research
Center
Athens U
Athens, Tech Univ

CERN Delphi, cont'd.

Bergen U
INFN, Bologna
Brussels U, Iihe
CERN
Bohr Inst
Cracow, Inp-Exp
Dubna, Jlnr
Genoa U And INFN, Genoa
Grenoble U
Helsinki U
Iowa State U
Karlsruhe U, Iekp
Lisbon, Lifep
Liverpool U
Ljubljana U
Lund U
Lyon, Ipn
Marseille, Cppm
Milan U And INFN, Milan
Mons U
Orsay, Lal
Oslo U
Oxford U
Padua U And INFN, Padua
College De France
Paris, Curie Univ Vi, Lpnhe
Charles U
Prague, Inst Phys
Rio De Janeiro U
Rio De Janeiro, Pont Univ
Catolica
Rio De Janeiro, Cbpf
Rome, ISS And INFN, Rome
Rome U, Torvergata And INFN, Rome
Rutherford
Saclay
Santander U
Serpukhov, IHEP (Russia)
Stockholm U
Strasbourg, CRN
Trieste U And INFN, Trieste
Turin U And INFN, Turin
Udine U And INFN, Udine
Uppsala U
Valencia U
Vienna, Oaw
Warsaw, Inr
Wuppertal U
Lancaster U
Madrid U
Oviedo U
Aachen, Tech Hochsch, I Phys
Inst And Aachen, Tech
Hochsch, III Phys Inst
NIKHEF, Amsterdam
Michigan U
Annecy, Lapp
Johns Hopkins U
Beijing, Inst High Energy Phys

CERN Delphi, cont'd.

Humboldt U, Berlin
 INFN, Bologna
 Tata Inst
 Boston U
 Northeastern U
 Bucharest U
 Budapest, CRIP
 Harvard U
 MIT
 Florence U And INFN, Florence
 CERN
 World Lab, Geneva
 Geneva U
 Hefei, Cust (China)
 Helsinki U
 Lausanne U
 Los Alamos
 Lyon, IPN
 Madrid, Ciemat
 INFN, Milan
 Moscow, Itep
 Naples U, Ifs And INFN, Naples
 Cyprus U
 Nijmegen U And Nikhef, Nijmegen
 Oak Ridge
 Cal Tech
 Perugia U And INFN, Perugia
 Carnegie Mellon U
 Princeton U
 Rome U And INFN, Rome
 St Petersburg, Inp
 UC, San Diego
 Santiago De Compostela U
 Sofiya, Automation Sci
 Instrum Lab
 Korea Inst Sci, Taejon
 Alabama U
 Purdue U
 PSI, Villigen
 Desy-Ifh, Zeuthen
 Zurich, ETH
 Hamburg U
 Taiwan, HEP Group

CERN Opal

Aachen, Tech Hochsch, III
 Phys Inst
 Alberta U
 Birmingham U
 Bologna U And INFN, Bologna
 Bonn U
 Cambridge U, Cavendish Lab
 CERN
 Chicago U
 Duke U
 Freiburg U
 DESY
 Heidelberg U, Ihep

CERN Opal, cont'd.

Indiana U
 Kobe U (Japan)
 Birkbeck Coll
 Brunel U (UK)
 Queen Mary - Westfield Coll
 University Coll, London
 Manchester U
 Maryland U
 Montreal U
 Oregon U
 Carleton U
 CRPP, Ottawa
 UC, Riverside
 Rutherford
 Saclay
 Technion
 Tel Aviv U
 Tokyo U
 British Columbia U
 Victoria U
 Weizmann Inst

Cornell-CLEO

Cal Tech
 UC, San Diego
 UC, Santa Barbara
 Carleton U
 McGill U
 Colorado U
 Cornell U, Lns
 Florida U
 Harvard U
 Illinois U, Urbana
 Ithaca Coll
 Kansas U
 Minnesota U
 Suny, Albany
 Ohio State U
 Oklahoma U
 Purdue U
 Rochester U
 Southern Methodist U
 Syracuse U
 Vanderbilt U
 Virginia Tech

FNAL-831

UC, Davis
 Colorado U
 Fermilab
 Frascati
 Illinois U, Urbana
 Korea U
 Lebedev Inst
 INFN, Milan
 Milan U
 North Carolina U
 Notre Dame U

FNAL-831, cont'd.

Pavia U
Puerto Rico U, Mayaguez
South Carolina U
Tennessee U
Vanderbilt U
Wisconsin U, Madison

FNAL-CDF

Argonne
Bologna U And INFN, Bologna
Brandeis U
UCLA
Chicago U
Duke U
Fermilab
Frascati
Harvard U
Illinois U, Urbana
IPP, Canada And McGill U And
Toronto U
Johns Hopkins U
KEK, Tsukuba
LBL, Berkeley
MIT, LNS
Michigan U
Michigan State U
New Mexico U
Osaka City U
Padua U And INFN, Padua
Penn U
Pittsburgh U
INFN, Pisa And Pisa, Scuola
Normale Superiore And Pisa U
Purdue U
Rochester U
Rockefeller U
Rutgers U
Taiwan, Inst Phys
Texas A And M
Tsukuba U (Japan)
Tufts U
Wisconsin U, Madison
Yale U

FNAL-D0

Andes U, Bogota
Arizona U
Brookhaven
Brown U
UC, Davis
UC, Irvine
UC, Riverside
Rio De Janeiro, Cbpf
Mexico, Ipn
Columbia U
Delhi U
Fermilab
Florida State U
Hawaii U
Illinois U, Chicago

FNAL-D0, cont'd.

Indiana U
Iowa State U
Korea U
Pusan National U
Lbl, Berkeley
Maryland U
Michigan U
Michigan State U
Moscow State U
Nebraska U
New York U
Northeastern U
Northern Illinois U
Northwestern U
Notre Dame U
Oklahoma U
Panjab U
Serpukhov, Ifve
Purdue U
Rice U
Rochester U
Saclay (France)
Seoul National U
Sunny, Stony Brook
Tata Inst
Texas U, Arlington
Texas A And M

FNAL-KTEV

UCLA
UC, San Diego
Chicago U
Colorado U
Elmhurst Coll
Fermilab
Osaka U
Rice U
Rutgers U
Wisconsin U, Madison

DESY-ZEUS

Argonne
Bologna U And INFN, Bologna
Bonn U
Bristol U, Wills Lab
Brookhaven
Calabria U And INFN, Cosenza
Nevis Labs, Columbia U
Cracow, Inp-Exp
Cracow, Inst Phys Nucl Tech
Jagellonian U
DESY
Desy-Ifh, Zeuthen
Florence U And INFN, Florence
Frascati
Freiburg U
Glasgow U
Hamburg U
Imperial Coll, London
Iowa U

DESY-ZEUS, cont'd.

Julich, Forschungszentrum
 Korea U
 Louisiana State U
 Madrid, Autonoma U
 Manitoba U
 McGill U
 Moscow Phys Eng Inst
 Moscow State U
 NIKHEF, Amsterdam And
 Amsterdam U
 Ohio State U
 Oxford U, Npl
 Padua U And INFN, Padua
 Penn State U
 Rome U And INFN, Rome
 Rutherford
 UC, Santa Cruz
 Siegen U
 Tel Aviv U
 Tokyo U, Ins
 Tokyo Metropolitan U
 Turin U And INFN, Turin
 Turin U, Alessandria And
 INFN, Turin
 Toronto U
 University Coll, London
 Virginia Tech
 Warsaw U, Iep
 Warsaw, Inst Nucl Studies
 Weizmann Inst
 Wisconsin U, Madison
 York U, Canada

KEK-AMY

Rochester U
 South Carolina U
 Korea U
 Louisiana State U
 Beijing, Inst High Energy
 Phys
 Virginia Tech
 UC, Davis
 Hawaii U
 KEK, Tsukuba
 Gyeongang National U
 Konan U
 Minnesota U
 Niigata U
 Nihon Dental Coll
 Rutgers U
 Saga U, Japan
 Seoul National U
 Kyungpook National U
 Chuo U, Tokyo
 Saitama U

CERN-Nomad

Annecy, LAPP
 Massachusetts U, Amherst
 Calabria U
 CERN

CERN-Nomad, cont'd.

Dortmund U
 Dubna, JINR (Russia)
 Florence U And INFN, Florence
 Harvard U
 Johns Hopkins U
 Lausanne U
 Melbourne U
 Moscow, INR
 Padua U And INFN, Padua
 Paris, Curie Univ Vi, Lpnhe
 And Paris, Univ Vii, Lpnhe
 Pavia U And INFN, Pavia
 Pisa U And INFN, Pisa
 Saclay
 Sydney, Ansto
 Sydney U
 UCLA
 Boskovic Inst, Zagreb

SLAC-SLD

Adelphi U
 Boston U
 Brunel U
 Cincinnati U
 Colorado U
 Colorado State U
 Columbia U
 Ferrara U And INFN, Ferrara
 Frascati
 Illinois U, Urbana
 LBL, Berkeley
 Massachusetts U, Amherst
 MIT
 Nagoya U
 Oregon U
 Perugia U And INFN, Perugia
 Rutgers U
 Rutherford
 SLAC
 Sogang U
 Tennessee U
 Tohoku U
 UC, Santa Barbara
 UC, Santa Cruz
 Vanderbilt U
 Washington U, Seattle
 Wisconsin U, Madison
 Yale U

Gran Sasso-MACRO

Bari U And INFN, Bari
 Bologna U And INFN, Bologna
 Boston U
 Cal Tech
 Drexel U
 Frascati
 Gran Sasso
 Indiana U
 Aquila U
 Lecce U And INFN, Lecce
 Michigan U

Gran Sasso-MACRO, cont'd.

Naples U, Ifs And INFN, Naples
Pisa U And INFN, Pisa
Rome U
Texas A And M
Turin U And INFN, Turin
Bartol Research Inst
Sandia

Russia-SAGE

Moscow, INR
Los Alamos
Penn U
Louisiana State U
Princeton U

Sudbury Neutrino Observatory (SNO)

Chalk River, Aecl
Queens U, Kingston
CRPP, Ottawa

Sudbury Neutrino Observatory (SNO) , cont'd.

Carleton U
Guelph U
Laurentian U
British Columbia U
Penn U
Princeton U
Los Alamos
LBL, Berkeley
Oxford U

Underground - Soudan 2

Argonne
Minnesota U
Argonne And Minnesota U
Oxford U
Rutherford
Tufts U
Western Washington U

Nuclear Physics

*William A. Zajc
Nevis Laboratories
Columbia University*

The Nuclear Physics Program

The central objectives of the Nuclear Physics program are to understand the interactions, properties, and structures of nuclei and nuclear matter and to understand the fundamental forces of nature as manifested in atomic nuclei. While these goals have changed only slightly in the past fifty years, the methods and approaches by which these goals are pursued have changed dramatically over the past decade. Previously, there were many small facilities available, generally one per institution, and many Nuclear Physics researchers performed experiments using machines at their own sites. Most experiments in Nuclear Physics measured only a few parameters per event and involved only one to three institutions per experiment. As a result, Nuclear Physics research tended to be localized, small, and independent. Networks were not essential.

In the last ten years, in response to advances in theoretical descriptions, experimental techniques, and computer processing power, many Nuclear Physics investigators have moved to higher energies or chosen to perform much more complex experiments at low energies. This has led to a trend similar to that which occurred in the evolution of High

Energy Physics, that is, a reduction in the number of facilities combined with an increase in the size of collaborations.

Nuclear Physics now has active experiments that measure thousands of parameters and involve tens of participating institutions, and the experiments planned for the Relativistic Heavy Ion Collider (RHIC) will be roughly an order of magnitude larger. Nuclear Physics experiments have matched or exceeded High Energy Physics experiments in several measures of complexity. For instance, with 16,000 readout pads recording both time and pulse-height, the Equation-Of-State time-projection chamber built at LBL is comparable in complexity to the detectors used in many "small" High Energy Physics experiments. The same trend can be seen in the Nuclear Physics work being done with heavy ion beams at the Brookhaven Alternating Gradient Synchrotron (AGS), where particle multiplicities are already higher than those that are expected in High Energy Physics experiments at CERN's Large Hadron Collider.

This trend notwithstanding, important differences still exist between High Energy Physics and Nuclear Physics. The main differences that affect the need for network resources are location, diversity, and scale. The location issue simply reflects the fact that there are "pure" Nuclear Physics sites, such as the Continuous Electron Beam Accelerator Facility (see Figure 8-1) and the Sudbury Neutrino Observatory (SNO), as well as Nuclear Physics sites (such as BNL and LBL) where there are also major High Energy Physics programs. Diversity and scale are more complex issues that are coupled to some degree. Nuclear Physics, consisting of numerous small groups and some larger collaborations, actually has a greater diversity of programs than High Energy Physics. The smaller groups reflect the historical origins of Nuclear Physics research in small, university-based studies, while some of the larger Nuclear Physics efforts reflect the trend toward larger collaborations based at major research facilities. These larger Nuclear Physics collaborations have requirements very similar to those of large High Energy Physics programs.

Current Requirements

Most Nuclear Physics researchers now have access to a suite of basic ESnet services. Many research programs have come to rely in an essential way on the availability of these services, to the extent that severe disruptions in scientific progress would result should they be suspended. Therefore, these services now provide a working definition of current requirements. They include:

1. *Electronic mail*. Effectively instantaneous (less than 1 min.) transmission of simple text files.

Figure 8-1. The Continuous Electron Beam Accelerator Facility (CEBAF)

2. *File transfer capability.* It must be possible to send both binary and text files transparently (i.e., independent of internal representation) from node to node. Transmission rates of 500 kbps are required for existing Nuclear Physics experiments.
3. *X (virtual) terminal capability.* For X-terminals that provide access to a complete suite of network services, interactive response time should be less than 0.2 second for connections between two network nodes.
4. *Protocols.* TCP/IP-based networking has become predominant, as new collaborations have largely gravitated towards UNIX-based systems. However, continued support of full native-mode DECnet communication remains essential to operations and data analysis in existing experiments.
5. *Connectivity.* Connections are required to all laboratories, institutions, and computing centers in the world where U.S. nuclear physicists are involved. A large number of nuclear physicists are supported by NSF, so the current direct peering of ESnet to NSFnet is essential.
6. *Security.* The tremendous productivity gains made possible by the current ESnet must not come at the price of reduced security for local systems.
7. *Availability.* The Nuclear Physics program is highly dependent on a stable network and will become even more dependent on such a facility in the future. The following requirements are essential and reasonable:
 - Greater than 99% uptime for connectivity between any two nodes
 - A mean time between failures greater than one month
 - A mean time to repair that is typically less than two hours
 - Downtimes of greater than half a day should not occur more than once per year.

Near-Future Requirements

The following media have recently emerged as powerful new applications or methods with implications for additional ESnet usage in the near future. All are presently in use, although not with the near-universal pervasiveness of those listed in the previous section.

1. *World Wide Web.* The WWW has provided a new paradigm for accessing, presenting, searching and archiving information across the network. The Web has proven a very powerful tool for organizing information, particularly when accessed via Mosaic. Some existing Nuclear Physics experiments are moving their existing

files to this system, and some future experiments have adopted it as their primary archival system.

2. *Distributed computing environments* (DCEs). Larger collaborations inevitably lead to distributed resources for computing. In the future, the development and maintenance of distributed applications will be essential to the progress of planned Nuclear Physics experiments and, to some extent, to collaborations in nuclear theory. CEBAF has taken a leading role in identifying DCE-related issues common to many Nuclear Physics programs.
3. *Desktop videoconferencing*. The use of ESnet to support multipoint desktop videoconferencing has recently been demonstrated. This development, coupled with the trend toward large and distributed collaborations, will inevitably lead to greater utilization of ESnet resources, especially in view of the very modest hardware investment required on the part of users who wish to exploit these resources.

Future Requirements

Two major Nuclear Physics facilities with significant implications for ESnet utilization will begin operation before the end of the decade. The first, the Continuous Electron Beam Accelerator Facility, will study inelastic electron-nucleus scattering. The second machine, the Relativistic Heavy Ion Collider, will study high-energy nucleus-nucleus collisions.

CEBAF

Beginning in FY95, CEBAF experiments will begin taking data at recording rates of 10–100 gigabytes per day. In FY97, this rate will rise to approximately one terabyte per day. Initially (i.e., in FY95), the processing power required for data reduction will be 2000 MIPS and will increase by a factor of ten by FY97. This processing power will be provided by a local computer farm that will be accessible to off-site users via ESnet. The on-line storage accessed by such users is currently 40 gigabytes per year and is expected to grow to 30–50 terabytes per year in FY97. By the end of FY98, CEBAF plans to have more than 100 terabytes of data on-line at any one time, with more than 500 gigabytes per month going onto ESnet. These data will be reduced by resources that will exceed 20,000 MIPS of computing power.

CEBAF currently has 437 authors on approved experiments. These authors are from 101 institutions in 18 countries. Network services will be critical in enabling CEBAF collaborators to perform experiments, analyze data, and publish results.

RHIC

RHIC operations will begin in FY99. Two large and at least two small experiments will take data at aggregate recording rates in excess of one terabyte per day. A significant fraction of the processing capacity needed for data reduction will be provided by the RHIC computing center, which will have 100 gigaflops of computing power and will be accessible to off-site users via ESnet. The on-line storage accessed by such users is expected to be on the order of 200 terabytes.

RHIC currently has 613 authors on approved experiments. These authors are from 70 institutions in 13 countries. With such widely distributed participants, the RHIC collaboration is fully exploiting the high-bandwidth connectivity currently provided by ESnet. This connectivity supports the distributed computing environment for the Monte Carlo simulations required to plan these experiments. ESnet also provides desktop videoconferencing, which makes it possible to stay on top of critical engineering and software development issues. The continued growth of these services is essential if RHIC collaborations are to perform their experiments, analyze their data, and publish their results in a timely fashion.

Fusion Energy

*Martin Greenwald
Plasma Fusion Center
Massachusetts Institute of Technology*

The mission of the Fusion Energy Program is to develop nuclear fusion as an environmentally attractive, commercially viable, sustainable energy source. To achieve this goal, it has been necessary to develop the science of plasma physics, a field closely related to 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. As a result, plasma theory has become computationally intensive. The evaluation of fusion research requires the acquisition and analysis of very large numbers of signals that describe plasma parameters as functions of either space and time or energy and time. Some types of analysis involve the generation of video images from these parameters, and this activity is also computationally intensive.

In this country, FE work is distributed across nearly 100 sites, with concentrations at about a dozen national laboratories, universities, and private contractors. Research on a similar scale is carried out in Japan, Western Europe, and Russia. However, both nationally and internationally, progress in FE research is linked to a handful of large experiments that depend on facilities that are too expensive to duplicate. As these experiments increase in size, we see a concurrent growth in the number and importance of collaborations between the large groups at the experimental sites and smaller groups located at universities and the

national labs. Future collaborations will be national (e.g., TPX, the Tokamak Physics eXperiment)¹ or international (e.g., ITER, the International Thermonuclear Experimental Reactor)² in scope. Collaborators in these experiments are spread over the entire FE community, and design efforts by national and international teams (as appropriate) are already under way.

Computing and Networking in FE Research

Supercomputer Use

Many phases of activity within the FE program generate problems whose solutions require the use of supercomputers. To cite just a few examples:

- Large-scale engineering codes are applied to numerous aspects of the design of magnetic confinement devices. These codes make it possible to analyze thermal, structural, neutronic, electromagnetic, and power parameters as well as overall system design.
- Transport codes are composed of many modules, each containing a good deal of physics, and can result in lengthy calculations.
- Three-dimensional equilibrium computations are needed for modeling stellarator processes and for studying toroidal ripple in tokamaks.
- Modeling the wide range of instabilities in a fusion device requires three-dimensional fluid and kinetic calculations with very fine resolution.
- At present, one of the major Grand Challenge problems in the FE program is the investigation of plasma turbulence and heat transport. The design of new machines and proposed reactors depends critically on our predictions of the magnitude of heat transport. To date, the fusion community has relied mainly on empirical scaling laws to predict heat transport. More accurate predictions are a widely recognized goal within the community.

Increasingly, high-speed communications are essential to the solution of such computational problems. The past sixteen years have seen a substantial increase in both the computing power available at NERSC and the scope of the calculations performed there. A single state-of-the-art turbulence calculation can now generate several gigabytes of data that describe the evolution in space and time of the various physical quantities of interest. Traditionally, this data was archived and analyzed at NERSC, and a minimum of data was transferred to remote sites for printing and plotting. Now, with the advent of affordable local mass-

storage devices (such as optical jukeboxes) and powerful desktop workstations, remote users can perform efficient in-depth analysis and graphical display of data at their home sites. The increase in such activity at users' home sites will create greater requirements for data communications between NERSC and remote sites. The demand for "seamless" integration between NERSC and systems at remote sites is expected to place even greater pressure on network bandwidth and availability.

Similar increases in demand for bandwidth and availability occurred some two years ago, following the creation of two new computer centers that house leading-edge massively parallel supercomputers. These two centers are LANL's Advanced Computing Laboratory and ORNL's Center for Computational Sciences. Officially designated as High Performance Computing Research Centers, they were established as part of the High Performance Computing and Communications Initiative. At this time, their full exploitation by the FE program requires increased network access to them for fusion researchers at remote sites who are participating in the Numerical Tokamak Grand Challenge.

Experimental Collaborations

Large fusion experiments produce in excess of 200 gigabytes of raw data per year. For experiments like TPX or ITER, this number could be larger by a factor of 10 or more. Considering the current trend toward transferring data to remote sites for analysis, the TPX and ITER programs will increase the FE program's data communications requirements considerably.

Experimental collaborations generate two kinds of requirements. One requirement is for network access to experimental data so that researchers can remotely analyze and display it. The other requirement involves remote-control-room applications, in which experiments are controlled and monitored from remote sites over the network. The requirement for access to experimental data has been around for many years, with the rate of data transfer and processing power increasing dramatically from year to year. By contrast, we are just beginning to gain experience with remote-control-room applications.

A significant step in the latter direction was taken in connection with PPPL's Tokamak Fusion Test Reactor (TFTR) project. PPPL personnel developed a prototype remote-control diagnostic station that allowed TFTR collaborators located at the University of Wisconsin (UW) to participate directly in the operation of the beam-emission spectroscopy diagnostic housed at PPPL. This remote-control facility provided pseudo-real-time display of data analysis, shot clock monitoring, and full remote control of the diagnostic. The prototype was successful, the only

drawback being sluggishness due to inadequate network bandwidth between PPPL and UW.

As a result of this success, other remote TFTR collaborators have considered using the same environment for their diagnostic work. These collaborators include researchers at MIT who are working on the Lithium Pellet Injector and researchers at ORNL who are working on the Oak Ridge Reflectometer. Both of these efforts have grown since the start of TFTR's D-T phase.

The expanded use of remote-control facilities will further increase demand for videoconferencing (both desktop and conference-room) to support daily run-planning meetings and allow "live" participation in TFTR control room activity by remote diagnosticians. Figure 9-1 illustrates the use of these videoconferencing facilities. That figure shows an on-site TFTR researcher holding a two-way desktop videoconference with the remote collaborator pictured in the lower monitor. The on-site researcher can simultaneously monitor control room activity in another window on the same screen. The upper monitor shown in Figure 9-1 displays various kinds of graphical information, which might include diagnostic data or results from the previous experimental shot. Like the on-site researcher, the remote videoconference participant has access to the control room video as well as the graphical displays. Using these shared network facilities, researchers in widely separated locations can evaluate previous results and other pertinent information while monitoring control room developments. They can then factor in these evaluations and observations as they set the parameters for the next shot.

A different remote-collaboration capability has been implemented by M-Division researchers from LLNL who are who are associated with the DIII-D tokamak at General Atomics. These collaborators were able to employ a well-established computer system at LLNL for remote analysis of DIII-D experimental data, which is shipped back to GA after it is analyzed. Remote collaborations on the recently commissioned C-Mod tokamak at MIT are also beginning. Such collaborations should greatly increase in scope over the next few years, as staff people at PPPL and ORNL become available to participate in experiments at remote sites.

Theoretical Collaborations

Theoretical collaborations began to make use of distributed workgroup computing approximately two years ago. Following the introduction of this mode of collaborating, ESnet-based traffic has doubled each year. The input/output traffic generated by NERSC-based production has significantly increased during this period, due to the use of distributed file systems located locally and accessed remotely via the

Figure 9-1. Videoconferencing facilities in use by TFTR collaborators

UNICOS-based Cray supercomputers. The dramatic increase in NERSC traffic has resulted partly from the fact that these distributed file systems allow for the local preparation of source code and input data sets that are used directly by the UNICOS compilers. (These distributed file systems allow production to be done in the traditional way, with data transferred to the Crays at NERSC and processed on those systems through the application of NERSC-based production codes.) On the output side, the production data sets produced by the supercomputers are often written directly to remote file systems.

This geometric growth in network traffic has already resulted in periods of saturation during prime time for ESnet access and is expected to continue into the future, as more sophisticated distributed applications—especially those that involve scientific visualization—are brought on line.

Design Collaborations

During the design phase for the next generation of fusion devices (i.e., TPX and ITER), large-scale engineering analysis will absorb a significant share of the available computing resources. (As noted above, the engineering codes allow the analysis of electromagnetic, structural, neutronic, and thermal factors.) Since the design groups are widely distributed, both nationally and internationally, the ability to communicate models and results of simulations and analyses will be an important requirement for ESnet.

The ITER Engineering Design Activity (EDA) is now underway at the three Joint Central Team (JCT) sites, San Diego, Naka, and Garching. The U.S. has also established a Home Team tasked with performing the research and development work needed for ITER support. The JCT is responsible for providing configuration control for the ITER EDA. This function is being performed by the ITER Integrated Process Management System (IPMS). IPMS is a computer-based application that facilitates the control of both technical and management information throughout the ITER project. The database responsible for this system is ultimately expected to describe roughly 50,000 components and to grow in size to several hundred gigabytes. Access to this database will be required by users at the three JCT sites as well as by the Home Team members at their various sites. To facilitate timely access to the current database, a copy of it is updated daily at each of the JCT sites. ESnet is being used as the common link for both access and updates to this database. Other ITER networking requirements include such standard items as e-mail, file transfer facilities, and remote computer access. In addition, desktop videoconferencing will be required in the near future.

Current Networking Requirements

ESnet provides the communications infrastructure that ties the individual components of fusion research into a well-integrated international program. This is particularly important because fusion research is moving into a phase in which experimental collaborations diminish in number while growing in size and theoretical collaborations grow dramatically in both scope and scale.

Connectivity Requirements

Worldwide sites involved in FE research include some 100 in the US, 65 in Europe, 45 in Japan, and several in South America, China, and Australia. The principal sites are:

In the USA

Argonne National Laboratory
Columbia University
General Atomics*
Lawrence Berkeley Laboratory
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Massachusetts Institute of Technology*
National Energy Research Super Computer 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*

In Canada

Varenes: Hydro Quebec CCFM*

In France

Cadarache: CEN*

In Germany

Garching: IPP, Max Planck Institute*
Jülich: KFA*
Kuhlsruhe: KFK

In Italy

Frascati: CREF (ENEA)*
Padua: IGI*

In Japan

Kyoto: Kyoto Univ., Uji
Nagoya/Toki: NIFS*
Naka: JAERI*
Tokai: JAERI*

In Russia

Moscow: Kurchatov, Keldish Institutes*
St. Petersburg: Ioffe, Efremov Institutes

In Switzerland

Lausanne: CRPP, EPFL*

In the UK

Culham/Harwell: JET, Culham Laboratory*

* indicates the site of a major FE experiment.

For U.S. FE research sites, the highest priority for network connectivity has traditionally been access to NERSC. As theoretical and experimental collaborations have grown in size and importance, first-rate lab-to-lab connectivity has also become essential. In addition, because of the international nature of FE research, the importance of international links rivals that of the domestic connections.

Specific connectivity requirements are:

- Access to NERSC

- Support for experimentalists at remote sites who are operating diagnostics for and participating in experiments such as TFTR at PPPL, DIII-D at GA, C-Mod at MIT, TEXT at UTA, JET and START at Culham, ASDEX-UG at IPP Garching, TORE SUPRA at Cadarache, TEXTOR at Julich, FTU at Frascati, RFX at Padua, and JT-60U at Tokai
- Support for ITER design teams at San Diego, Garching, and Naka
- Support for the ITER research and development program³
- Communications support for the TPX design team, whose members are located at PPPL, ORNL, MIT and GA
- Support for physics and engineering research and development efforts dedicated to TPX⁴
- Gateways to the rest of the Internet, NASA's NSI, and the other federal backbone networks.

Functional Requirements

At a bare minimum, the FE program requires the standard network services, including electronic mail, file transfer, remote log-in, remote printing and plotting, and support for bulletin boards. To ensure support for such capabilities as full-screen editing, visualization, and windowing, it is preferable that the connections be made via high-bandwidth terrestrial lines with minimal latency. TCP/IP and DECnet are widely used in the fusion community, so multiple protocols should be supported. Lastly, there is a need for a network directory service that would support electronic white and yellow pages.

Additional services are needed to support the growing number of remote collaborations. Computing, visualization, and file management for a given task are already shared over the network. A large number of FE researchers are using a variety of distributed network services via workstations and windowing terminals. This style of computing requires not only greater network bandwidth and very high availability but also a new order of functionality to support what is often referred to as "seamless" distributed computing. Such functionality would include remote job entry, remote procedure calls, network file service, distributed editors, distributed databases, and so forth. In addition, FE researchers will soon need such capabilities as remote or distributed code management and distributed system management.

The nature of the FE program requires that these services be available on international as well as national networks. Because current design activities and experimental and theoretical collaborations are increasingly multi-institutional as well as multinational in scope, the FE community will require more powerful tools for interpersonal communications, such as videoconferencing and multimedia mail.

It is now assumed that the TPX and ITER programs will set up distributed control rooms to support data acquisition and instrument control for experimenters located at remote sites. The nature of the experimental methods employed in FE research requires close communication between all members of the physics teams responsible for acquiring and analyzing experimental data. While it is to be expected that many members of each lab's physics team will be present at the experimental site for a significant number of runs, it is unlikely that all of them will be present every day. Hence the proposal to link the control room to collaborators' home sites with data, voice, and video communications.

Performance Requirements

Today's ESnet, with its TCP/IP and DECnet components, its T1 backbone, and its gateways to the Internet, generally meets the current needs of the FE program. The most notable unmet need is for high-speed tail circuits to non-backbone sites. To meet present FE requirements, most of these tail circuits need to be upgraded, with 10 or so going to full T1 level. It should also be noted that as ESnet's traffic has grown, users have begun to encounter performance limitations. Clearly, deployment of T3 upgrades is coming none too soon. The processes described above (visualization, distributed computing, remote control of experiments, etc.) entail moving complex computing applications from a local-area network (LAN) into the WAN environment, and these procedures will require ever-increasing bandwidth. In this connection, it should be noted that at many sites the 10-Mbps Ethernets have become saturated, requiring migration to FDDI networks operating at 100 Mbps.

Forecast of Future Requirements

It is difficult to define future requirements with any precision, as users are generally not willing to ponder the effects of future applications that will change their technology and their working procedures in major ways. That being the case, we can best estimate network requirements by examining current and future activities that will require support.

Operational Trends in Current Experiments

Two operational trends on the immediate horizon have clear implications for network support requirements. The first of these is the transition to the use of network-based applications that will enable remote researchers to conduct operations, as nearly as possible, as though they were on-site. This trend will entail the implementation of remote-control environments and a sizable increase in the use of videocon-

ferencing. The implications for network-support requirements are fairly obvious, as these are both bandwidth-intensive applications, and the implementation of remote-control-room facilities will involve a considerable amount of developmental work.

The DIII-D project provides an example of a collaboration in which this operational transition is already under way. As noted above, DIII-D collaborators have already implemented remote analysis (i.e., at LLNL) of data from test shots done at GA in San Diego, CA. With improvements in the network link from LLNL to GA and LAN upgrades at both ends, we could see a remote-control-room environment connecting the two sites. The data/computations connections are fully in place, and there is little difference between the two sites in terms of data flow. With a remote-control-room environment established, the missing ingredient for remote collaborators would be the interactive discussions and weekly meetings that are missed if one is not physically present at the experimental site. If videoconferencing—both desktop and conference-room—were also incorporated into everyday operations, this lack of interactive access could be largely overcome. Providing these videoconferencing facilities would also increase the productivity of the collaborators and minimize their need to travel. An additional capability needed by the LLNL-based scientists is one-way video/audio monitoring of the control room at the experimental site. This capability is essential so that the remote experimenter can get a feel for the tempo of activity during operations.

The other important operational trend is the reassignment—in growing numbers—of researchers who have been working on local projects to duties that involve remote collaboration. For example, TFTR is scheduled to cease operations at the end of December 1994, leaving some of the staff at PPPL free to collaborate on experiments being conducted at other sites. Thus, between the end of 1994 and the beginning of TPX operations in 2000, PPPL expects to significantly increase participation in other fusion experiments worldwide, including DIII-D, ATF, C-Mod, JT-60, JET, TORE-SUPRA, and ASDEX. Although a few PPPL staffers will move and/or commute to each of these sites, the majority of person-years devoted to these collaborations by PPPL staffers is expected to be spent at PPPL. This operational approach necessitates good network connectivity in support of such functions as data access, remote diagnostic control and monitoring, and desktop videoconferencing. A PPPL-based remote-control-room facility is also being considered. ORNL will go through the same change in operational style after the definitive shutdown of ATF in 1995, with on-site ORNL staffers increasingly dedicated to remote collaborations as their ATF responsibilities end.

The Numerical Tokamak Project

The Numerical Tokamak Project has a very high priority within the Office of Fusion Energy because it addresses the anomalous transport problem, a high-leverage problem in the study of fusion energy. In addition, the Numerical Tokamak Project is one of six Computational Grand Challenges selected by the High Performance Computing and Communications Initiative. All major fusion laboratories and universities, as well as groups at JPL/Caltech and Cornell, are active participants in this project.

One of the goals of the gyrofluid simulations is to be able to do 3-D simulations with a $512 \times 512 \times 64$ grid. Such a simulation would calculate values for 4–6 fields (density, parallel speed, etc.) at each grid point, so it would generate over 800 Mbytes of data at each time step. To save this much data from one time step (in order to be able to restart at a later time) would require approximately four hours of transfer time across the current ESnet facilities. This entire process would normally be done at the Advanced Computing Laboratory (ACL) at Los Alamos, as that facility has a high-bandwidth link between its Connection Machine and its data storage facilities. However, at times it would be desirable to download the 3-D data from a series of time steps in order to do local-mode visualization and analysis on remote workstations. If this transfer were done via the current ESnet facilities, downloading even 1 byte per field per grid point would require approximately 32 minutes per time step. Finally, it would also be desirable to be able to produce time-dependent movies for purposes of data interpretation and visualization. Producing such movies would involve even more data and require even faster communications facilities than the data transfers described above.

Future Experiments

The Tokamak Physics Experiment

The Tokamak Physics Experiment is a national project involving approximately seventeen institutions. Though the experiment will be located at PPPL, a number of fusion laboratories will have major roles in the design of TPX and in the later industrial procurement process. The goal is to generate the first plasma within seven years.

As a major distributed DOE design effort, this project will be a good test bed for advanced communications technology. Two capabilities are essential: (1) to be able to communicate design engineering information quickly and efficiently, and (2) to provide for physics planning and project management via videoconferencing. Current plans for videoconferencing include approximately sixty-four hours of

engineering/physics meetings per month and eight hours of project integration meetings per month.

A programmatic requirement for high-speed network links will unfold within the early stages of the TPX project, primarily in support of CADD (Computer Aided Design and Drafting) activity during the engineering/design phase. Two-dimensional drawings as well as three-dimensional models will be generated by designers using various CADD systems at PPPL, LLNL, ORNL, MIT and at various industrial sites. Timely, direct access to such information will greatly improve the integration process and minimize the engineering delays that result from using outdated information. CADD files vary in size from 400 kbytes to 8.4 Mbytes. Based on the number and size of the files to be transferred weekly by each participating site, we estimate that the file transmission time required via the current network will be approximately two hours per day per site.

By the time TPX is in operation, users at their home institutions will require transparent access to the entire archive of scientific data. The TPX physics 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 it is acquired. The successful implementation of such procedures would require each remote group to have excellent communications with the TPX computer system, the TPX machine operations staff, the TPX physics group at PPPL, and physics groups at other remote sites.

The International Thermonuclear Experimental Reactor (ITER)

The ITER Engineering Design Activity is fully under way within the scope of efforts defined to date, and extensive collaborations have begun between the selected international design centers. There is also a high level of activity among the U.S. ITER Home Team members. The major domestic network-related activity in this area is intensive Home Team interaction in the areas of R&D and project management. This activity has been earmarked as an application for videoconferencing in the future, and the use of videoconferencing for Home Team interactions will continue to grow as ITER EDA activity increases. The requirement to support this and other ITER videoconferencing applications will become part of the general ESnet requirements.

International Links

FE research has relied on international collaboration since its worldwide declassification in 1957. Such collaborations are supported by an extensive interchange of ideas and scientific personnel worldwide. In

addition, technology is readily exchanged between the countries involved, and components of experimental devices are fabricated all over the world. As ITER becomes the focus of the international program, such collaborations are becoming central to our work. In the future, we expect that international links will be as important to FE as our domestic connections are today.

Future Bandwidth Requirements

We can arrive at roughly equivalent estimates of future bandwidth requirements by two separate methods. The first begins with the assessment that all the major sites have local-area networks that provide 10–100 Mbps of bandwidth and use a fair fraction of this capacity, at least at peak-load periods. Supporting the applications and services described above will require equivalent functionality over the wide-area network. Summing loosely over the sites and assuming that communications between sites accounts for approximately 10% of the local traffic, we can estimate the WAN bandwidth requirements at 10–100 Mbps.

We can arrive at a similar estimate by noting the proliferation of workstations on the desks of FE physics and engineering staffers. These workstations now number in the hundreds and will likely reach several thousand before long. It has been estimated that when a workstation is actively using network resources, it requires between 0.1 Mbps and 1.0 Mbps of bandwidth. Based on these estimates alone, we can project future network bandwidth requirements by multiplying together the following factors:

- The number of workstations in our research community (= 1000)
- The utilization factor per workstation (= 0.1)
- The fraction of workstation utilization that represents wide-area network use (= 0.1)
- Network bandwidth used per workstation (= 0.3 Mbps)
- The load-factor “cushion” required to ensure that peak loads can be supported (= 10)
- The factor by which cpu speeds are expected to increase over the next five years (= between 3 and 10).

This approach suggests that by 1999 between 100 and 300 Mbps of network bandwidth will be required to support fusion energy research.

These two projections do not account for the bandwidth increases that will be required to support such applications as multimedia e-mail, videoconferencing, and distributed CAD. If we factor in these requirements, a very conservative estimate of future bandwidth requirements would be tens of Mbps in the very near future, rising toward 100 Mbps within a 5–10-year time frame.

Notes

1. W. M. Nevins, R. J. Goldston, D. N. Hill, S. C. Jardin, S. S. Medley, et al., "Mission and Design of the Tokamak Physics Experiment/Steady State Advanced Tokamak (TPX/SSAT)," presented at the 14th International Conference on Plasma Physics and Controlled Nuclear Fusion Research, Sept. 30–Oct. 7, 1992, Wurzburg, Germany (paper IAEA-CN-56/F-1-5); printed in *Nuclear Fusion Supplement*, **3**, 279 (1993).
2. K. Tomabechi, J.R. Gilleland, Yu. A. Sokolov, R. Toshci, and the ITER team, "ITER Concept Definition," *Nuclear Fusion*, **31** (6), 1135 (1991).
3. ITER Related Physics R&D Needs, ITER-IL-PH-16-0-18
4. "Research and Development Plan," TPX Project Memo No. 91-930708-PPPL/JSchmidt-01, July 1993.

Basic Energy Sciences

*Raymond A. Bair
Environmental Molecular Sciences Laboratory
Pacific Northwest Laboratory*

The Divisions of the Office of Basic Energy Sciences

The Office of Basic Energy Sciences (OBES) is divided into five divisions: Chemical Sciences, Materials Sciences, Engineering and Geosciences, Energy Biosciences, and Advanced Energy Projects. Each division funds basic, directed, and/or applied research in support of DOE's mission in energy research and environmental restoration. BES work is performed through major projects at DOE laboratories, a small number of projects at industrial laboratories, and hundreds of grants at U.S. universities. BES research interactions and collaborations extend to many other universities, both nationwide and worldwide.

Much of the research at DOE laboratories utilizes unique facilities, some of which are national user facilities. Major user facilities permit forefront research to be conducted in areas important to DOE by scientists from industry and universities as well as by DOE contractors/grantees. The OBES supports such specialized user facilities at national

laboratories. These facilities facilities could not be duplicated economically elsewhere, and they make it possible to conduct research that could not be performed by other means. These facilities are available to scientists from all DOE laboratories, from universities, and from industry, and they are typically used by all three communities.

The Division of Chemical Sciences

This division is divided into two research branches, Fundamental Interactions and Chemical Sciences Processes and Techniques. Each of these branches is further subdivided into programs focused on specific disciplines. The Fundamental Interactions Branch includes programs in Photochemical and Radiation Sciences, Chemical Physics, Atomic Physics, and Facilities Operations. The Photochemical and Radiation Sciences program emphasizes the exploration of fundamental photochemical processes that are involved in the capture and conversion of solar energy. The Chemical Physics program supports research into the fundamental molecular processes that relate to DOE's mission in the areas of combustion, catalysis, and environmental restoration. The Atomic Physics program seeks to obtain the most accurate and complete knowledge of the fundamental properties and interactions of photons, electrons, atoms, and ions.

The Chemical Sciences Processes and Techniques Branch includes research programs in Chemical Energy, Separations and Analysis, Heavy Element Chemistry, Chemical Engineering Sciences, and Advanced Battery Technology. The Chemical Energy program includes basic chemistry research into chemical transformations or conversions that are fundamental to new or existing concepts of energy production and storage. The Separations component of the Separations and Analysis program supports basic research directed toward improving our understanding of methods for separating mixtures of gases, liquids, solids, and their component molecules, cations, and anions. The same program's Analysis component supports research on phenomena basic to analytical methods. The Heavy Element Chemistry program focuses on the study of the chemical properties and certain of the physical properties of actinide elements. The Chemical Engineering Sciences program addresses the energy aspects of such chemically related engineering topics as thermodynamics, turbulence in combustion processes, and rates of physical and chemical processes. The Advanced Battery Research and Development Program supports research to develop new generic battery technologies oriented toward the nonautomotive consumer market.

User facilities operated by the Division of Chemical Sciences include:

- The Combustion Research Facility (CRF) at Sandia, Livermore

- The High Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory
- The Stanford Synchrotron Radiation Laboratory (SSRL) at Stanford University
- The National Synchrotron Light Source (NSLS) at Brookhaven (shared with the Division of Materials Sciences)

In addition, the Division of Chemical Sciences operates the Radiochemical Engineering Development Center at Oak Ridge National Laboratory as a service to the scientific community.

The Division of Materials Sciences

The Division of Materials Sciences supports basic research on materials-related phenomena that are important to all energy systems. The aim of this research is to provide the base of materials knowledge required to advance the nation's energy programs. Materials science research is traditionally characterized as small science with large numbers of individual investigators, but larger collaborations and the use of remote facilities play an increasingly important part in this research.

Materials science research activities fall mainly into four categories: metallurgy and ceramics, solid state physics, materials chemistry, and facility operations. Metallurgy and ceramics efforts focus on the structure of materials, mechanical and physical properties, radiation effects, and engineering materials. Solid state physics efforts include both experimental and theoretical research, neutron scattering, particle-solid interactions, and engineering physics studies. Materials chemistry research includes synthesis and chemical structure, polymer and engineering chemistry, and high-temperature and surface chemistry.

User facilities operated by the Division of Materials Sciences include:

- The Intense Pulsed Neutron Source (IPNS) at Argonne National Laboratory
- The High Flux Beam Reactor (HFBR) at Brookhaven National Laboratory
- Neutron Scattering at the High Flux Isotope Reactor at Oak Ridge National Laboratory
- The Materials Program at the Stanford Synchrotron Radiation Laboratory at Stanford University
- The National Synchrotron Light Source at Brookhaven (shared with the Division of Chemical Sciences)
- The Advanced Light Source (ALS) at Lawrence Berkeley Laboratory (see Figure 10-1)
- The Advanced Photon Source (APS) under construction at Argonne National Laboratory

Figure 10-1. LBL's Advanced Light Source (domed structure in center)

- The Materials Preparation Center at Ames Laboratory
- The Electron Microscopy Center for Materials Research at Argonne National Laboratory
- The Center for Microanalysis of Materials at the University of Illinois, Urbana-Champaign
- The National Center for Electron Microscopy at Lawrence Berkeley Laboratory
- The Shared Research Equipment (SHaRE) Program at Oak Ridge National Laboratory
- The Surface Modification and Characterization Research Center at Oak Ridge National Laboratory
- The DOE Center of Excellence for the Synthesis and Processing of Advanced Materials at Sandia National Laboratories, Albuquerque, NM
- The Materials Program at the Combustion Research Facility at Sandia National Laboratories, Livermore, CA

The Division of Engineering and Geosciences

The Division of Engineering and Geosciences conducts two research programs, Geosciences Research and Engineering Research. The mission of the Geosciences Research program is to provide the solid scientific understanding needed to develop efficient, economical, and environmentally sound energy technologies for the future. Fulfilling this mission requires fundamental knowledge of the processes involved in the consumption of energy and natural mineral resources and the treatment of resulting byproducts. Research efforts span the fields of geophysics and earth dynamics, geochemistry, energy resource evaluation and utilization, hydrogeology and exochemisrtry, and solar-terrestrial interactions.

The Engineering Research Program carries out fundamental research in support of related technology programs within DOE. The goal of this research is to extend the knowledge of engineering practices that will aid in the development of energy-saving industrial processes. Special emphasis is placed on projects which, if successfully concluded, will benefit more than one energy technology. Present efforts include research programs in:

- Mechanical Sciences, including fluid mechanics, heat transfer and solid mechanics
- System Sciences, including studies of process control and instrumentation

- Engineering Analysis, including nonlinear dynamics, databases for thermophysical properties of fluids, and modeling of combustion processes and bioprocessing.

The Division of Engineering and Geosciences operates the Center for Engineering Systems Advanced Research (CESAR) at Oak Ridge National Laboratory as a collaborative research facility.

The Division of Energy Biosciences

The Energy Biosciences program studies plants and microorganisms in order to investigate fundamental mechanisms that are important in the design of new energy-related biotechnologies. This work focuses on understanding the role of plants and microbes as renewable resources for fuel (hence as fossil fuel substitutes), as agents of environmental restoration, and as agents in reactions that provide energy savings compared with analogous chemical reactions. Primary research into biological energy conversion is concerned with plant and microbial photosynthesis as well as with plant development due to interactions with the environment. Studies of the bioconversion of products examine:

- The metabolic pathways in plants that are important in the synthesis of fuels or chemicals
- The genetic and biochemical regulatory mechanisms in plants
- The synthesis, structure, function, and degradation of plant cell walls
- Symbiotic energy exchanges between plants and microbes.

A key objective of this research is laying a foundation for the development of future biotechnologies, and to this end the program includes work on genetic expression and the development of critical databases and instrumentation.

The Division of Advanced Energy Projects

Projects in the Division of Advanced Energy Projects explore the feasibility of applying energy-related concepts that derive from basic research. This program serves as a bridge between basic research and applied development programs, providing a mechanism for applying the results of basic research in ways that could eventually improve the nation's energy economy.

Future Requirements

BES programs address a wide range of topics in the areas of research, development, and facility-operations, with a focus on the energy

and environmental issues that are critical to DOE and the nation. This programmatic diversity is reflected in the range of computational and network resources required by BES investigators and in the use of ESnet facilities and services by these investigators.

This diversity notwithstanding, we can identify several trends in the use of computational resources and network services in BES activities. Computation, information processing, and electronic communications are playing ever larger roles in scientific endeavors, and this is especially the case in fundamental research programs. In a given subject area, OBES-funded investigators may be widely distributed geographically, and the wider community of non-OBES-funded researchers with whom they interact is likely to be distributed worldwide. BES research is carried out in many DOE facilities, including all of the multiprogram DOE laboratories (ANL, BNL, INEL, LBL, LLNL, LANL, ORNL, PNL, and SNL), in single-program DOE labs like Ames Laboratory and the National Renewable Energy Laboratory, in mission-specific DOE labs like the Stanford Synchrotron Radiation Laboratory, and at many U.S. universities. OBES-funded researchers at any of these sites may employ experimental or computational resources at multiple remote sites. For example, remote control of experiments may be implemented at NSLS, and remote computations may be performed on the data from those experiments at the National Energy Research Supercomputer Center at LLNL, the Advanced Computer Research Facility at ANL, or the High Performance Computing Research Centers at LANL and ORNL.

ESnet and its network peers therefore have a critical role to play in ensuring the timely exchange of scientific ideas and results and the efficient use of computational and experimental resources. This distributed access to information and resources is becoming increasingly important to the success of BES research missions. By enabling people to work together more closely and by reducing duplication of effort and facilities, ESnet's services enable DOE to perform its mission more effectively. By making it easy for scientists to communicate, networks also support cross-fertilization between research disciplines. Finally, these same network capabilities provide a key infrastructure supporting the transfer of technology from DOE to U.S. industry.

Connectivity Requirements

Based on these communication trends, transparent network access is a prime requirement for BES usage of ESnet. BES programs require the interconnection of all the DOE labs. In addition, general connectivity to the U.S. university community is required to support the many BES university research programs. This connectivity is also needed to support DOE's increasing involvement in outreach to educational institutions.

Almost without exception, this connectivity exists today; maintaining it is a key requirement for the future. As U.S. industry and other federal agencies connect to various networks, there is good reason to provide strategic interconnection between those networks and DOE networks. Such links increase the accessibility of scientific expertise and that of the scientific data available on-line, which is rapidly growing in volume.

Overseas connectivity requirements for BES are still growing, and although most needed sites are reachable (e.g., European universities), the network service is sometimes poor or unreliable. Through its support of transatlantic and transpacific links, DOE is in a position to influence the interconnection of U.S. and evolving foreign networks to the benefit of DOE programs like BES. Overall, considering the amount of leverage to be gained from increased scientific collaboration among researchers both within and beyond DOE, the best strategy would be to provide transparent connectivity to all U.S. research institutions and, through strategic international connections, to all research institutions worldwide.

Bandwidth Requirements

As the importance of computing and information systems in DOE research increases, there is a concomitant increase in the requirement for network bandwidth. Bandwidth needs are driven by changes in electronics technology (e.g., faster computers and storage devices) and by new research capabilities (e.g., distributed computing, remote visualization, and videoconferencing). As noted above, ubiquitous connectivity is needed in the short term to support BES research. By contrast, BES short-term requirements for high bandwidth are limited to certain key sites. Over the longer term, high-bandwidth communication will probably become a ubiquitous requirement. Future network usage is expected to follow the pattern set in the past, when significant bandwidth increases opened up important new opportunities for researchers. Subsequent increases in network usage could not have been predicted by extrapolating demand curves that were based on the use of lower-bandwidth networks. As new network capabilities are disseminated and prove their value, they quickly become essential tools for scientists.

ESnet is crucial in supporting BES efforts that involve interlaboratory collaboration, shared lab resources and expertise, and emerging collaborative research environments (i.e., "laboratories without walls"). These new research paradigms employ network-intensive applications like remote visualization and experimentation, teleconferencing, shared whiteboards, and group document preparation. The use of many of these new and highly productive applications will require more than merely incremental increases in network bandwidth.

In the near term, BES programs would benefit from an aggressive and systematic deployment of increased network bandwidth between the DOE multiprogram laboratories, between these labs and other specific sites (including Ames Lab and Caltech), and between DOE network service providers and other high-bandwidth network providers, including those that are sponsored by other U.S. government agencies, those that are located at U.S. universities, and those that are in Europe. Greater network capacity is essential to ensure access for BES researchers to the computing facilities at ANL, LANL, NERSC, ORNL, SNL, and Caltech (and soon at PNL). The BES program is also planning for intersite transport of multigigabyte datasets, such as those generated by large-scale models or by the large data analyses performed by massively parallel computers. These plans clearly pose requirements that exceed ESnet's present capacity. Finally, the transport of visualization data will also require higher-bandwidth links between key DOE computing sites and specific sites within the general BES research community.

Requirements for Network Services

As the BES community expands its use of network facilities, it becomes increasingly important for ESnet to provide not only the requisite data transport capabilities but also DOE-wide network services. Two critical requirements for such network services are interoperability and ubiquitous availability, and ESnet has already taken several important steps toward meeting these requirements. At this juncture, further progress in this direction depends on the coordinated deployment of such network services as user authentication, video encapsulation, distributed computing tools, and information services. However, both network standards and information technologies are evolving rapidly. (Consider, for example, the surge in World Wide Web usage or the emerging Object Request Broker technology.) This highly fluid developmental environment makes it even more important for ESnet to coordinate the deployment of appropriate network services in concert with industry, universities, and other federal agencies.

The number of projects, meetings, and workshops conducted via ESnet facilities should continually increase for the foreseeable future. Network and distributed computing resources provided by ESnet should prove to be scalable toward larger implementations and compatible with future technologies. Finally, the network should be transparent to the user, regardless of that user's location or computer platform, and that transparency should include transparency to the distributed components of applications and services.

Metrics

The metrics required to support the successful use of ESnet and other network resources by BES researchers are much the same as those required in other leading-edge scientific programs. Network capabilities should anticipate and lead demand; to put it another way, the network should not be a factor that limits the effectiveness of BES research programs. ESnet should facilitate increases in productivity and help reduce the time required to solve BES research problems. When we find ourselves doing science in new and better ways, we will have achieved a key success.

Scientific Computing

*Sandy Merola
Information and Computing Sciences Division
Lawrence Berkeley Laboratory*

The Scientific Computing Program

The DOE Applied Mathematical Sciences (AMS) program is managed by OER's Office of Scientific Computing. The overall goal of the program is to improve DOE's ability to solve scientific and engineering problems that are critical to its mission and to the national interest by supporting research in and applications of advanced mathematical, computer, and computational sciences. In pursuit of this goal the AMS program performs two fundamental missions: (1) fostering the wide-ranging research in the mathematical, computational, and computer sciences needed to support and advance the physical and biological sciences, and (2) managing the Energy Research supercomputer access program, ER's high-performance computing research centers, and the Energy Sciences network.

The original AMS program was started in the early fifties at the instigation of John von Neumann, who saw the potential usefulness of the new digital computers in solving complex scientific and technological problems that were important to the Atomic Energy Commission (the predecessor of the DOE) and the nation as a whole. The AMS program

has changed considerably since that time but has remained consistent with its original purpose, which is to support mathematical and computational research that facilitates the use of the most advanced high-performance computer systems to enhance our understanding of science and technology. In 1983, the AMS program was incorporated into the newly formed Scientific Computing Staff, which reported to the Director of Energy Research. At that time, the responsibilities of the Scientific Computing Staff were divided into four areas: Analytical and Numerical Methods, Information Analysis Techniques, Advanced Computing Concepts, and Energy Sciences Advanced Computing. In 1993, the Scientific Computing Staff and its programs were folded into the newly formed Office of Scientific Computing, a program office of the Office of Energy Research. The creation of the Office of Scientific Computing reflected the recognition that improving the focus of computer and network-related activities was important both within the Office of Energy Research and beyond.

The components of today's AMS program are outgrowths of the division of responsibilities within the original Scientific Computing Staff. A comprehensive discussion of the current AMS program can be found in "The DOE Program in HPCC High-Performance Computing and Communications."¹ The responsibilities of the program's major components are summarized here.

1. *Basic Research and Human Resources.* This component supports basic research and human resources development activities in mathematics, computational sciences, and computer systems. As part of these efforts, this component supports educational activities and produces mathematical representations of physical systems as well as algorithms and techniques for predicting the behavior of physical systems on advanced computers.
2. *High Performance Computing Systems Research.* This component supports research and development on the software, computer languages, and hardware architectures of high-performance computing systems. Activities include the technological evaluation and characterization of new computer architectures and systems.
3. *Advanced Software Technology and Algorithms Research.* This component supports research and development on advanced software technology. A particular focus is the development of algorithms that enable effective application of advanced computers to scientific problems that are critical to the DOE. A prime goal is the early application of these technologies, algorithms, software tools environments, and computational techniques.
4. *Advanced Computation, Communications Research, and Associated Activities.* This component supports research, development, and operations needed to apply advanced computational,

communications, and mathematical techniques to scientific problems that are important to DOE and the nation. As one means of providing this support, this component ensures that OER-supported researchers have access to advanced computational and mathematical capabilities and resources. It also supports the associated high-performance computing research centers, access centers, communications infrastructures, and IRM functions required to ensure that this activity succeeds.

Requirements

AMS principal investigators play a wide range of roles in the development of advanced network technologies, from theoretical research to final implementation. At the same time, they are themselves leading-edge users of advanced network services. Therefore, their work both helps to structure and is structured by the requirements of modern network services. For example, AMS researchers generated much of the early demand for such "traditional" network services as remote log-in, the delivery of text-based e-mail, and the transfer of small to medium-sized files. These services are as thoroughly integrated into the AMS communications web as telephone or postal services, and AMS researchers are playing a leading part in fostering their use throughout the wider ER community.

In general, the requirements for network services are generated through a reciprocal process of user demand and "technological push." An example of "technological push" occurred recently, when greatly increased World Wide Web offerings were made available by the Mosaic network navigational tool. Even though there had been no stated user demand for such offerings, utilization of the network increased dramatically immediately after Mosaic became operational. The obvious utility of WWW and the resulting increase in user demand generated new requirements for networking bandwidth and services.

Other user demands have created additional requirements for either increased network bandwidth or improved network services—or both. For example, some applications that users have attempted to implement over the network have resulted in unsatisfactory response times. Users are clearly eager to make use of many of these applications over the network but are just as clearly unwilling to tolerate the associated network bottlenecks. In some cases, the demonstrated user demand remains unsatisfied.

The AMS work associated with high-performance computer systems reflects the need for access to unique facilities for computation, data access and manipulation, and data storage. The sophistication of modern methods typically requires communication and collaboration among an ever-growing number of ER scientific users. These researchers are

generally separated geographically from each other and from the computing facilities they must share. In this context of geographical separation, the growing importance and scale of scientific collaboration has created an urgent need for computer-supported cooperative work tools, such as desktop videoconferencing facilities and workstation-based shared design tools. AMS PIs have developed these tools and are providing an early distribution of them.

The Advanced Software Technology component of the AMS program aims to create a computational environment that serves the needs of some of the world's most advanced computing and computational researchers. Like the broader community of ER scientists, these collaborators are almost guaranteed to be geographically separated from each other and from the computing facilities they need. The computer network therefore becomes a critical component of the infrastructure required to advance research in this area. The network services necessary to support such an infrastructure are themselves only beginning to be understood and developed.

AMS efforts associated with the Advanced Computation and Communications Research component are dedicated to ensuring supercomputer access for Energy Research PIs, regardless of their locale. The efforts of this component also make ESnet and its associated services available to the entire ER community. These services are rapidly being extended to other programs within DOE via memos of understanding between the OER and those programs.

Current Network Utilization and Development by AMS Researchers

Because of their expertise in the mathematical and computer sciences and the nature of their contributions to those fields, the members of the AMS community are some of the world's leading users of computer networks. For example, AMS researchers were some of the earliest users of e-mail, file transfer, and remote access services. These services continue to be heavily used every day by AMS PIs. However, AMS scientists contribute to network use in a much more significant way through the demand for and development of new network services and technologies.

AMS scientists greatly accelerated the development of desktop (packetized) videoconferencing (DVC), which is now in the preproduction stage. Moving this product into full production is only one of the challenges currently being addressed by ESnet collaborators. These researchers have developed both the simple links required for one-on-one DVC and the Internet multicast topology required for multiuser DVC. This topology, which is known as the Mbone, is just beginning to be used outside the AMS community. Ultimately, the Mbone will be crucial in the development of computer-supported collaborative work environments.

Another network service that will soon be extended to a wider user base is multiple-site distributed computing, which involves the somewhat transparent use of computer resources located in geographically dispersed sites (as opposed to distributed computing within a single national lab or other site). The use of this service will spread from the AMS research environment into the broader ER community as AMS researchers improve the linkability of distributed processors (e.g., hypercubes) and as network bandwidth is increased to support such a distributed computing architecture.

Even as these new offerings become mature services, there are numerous other networking challenges for AMS researchers to surmount. For example, AMS PIs frequently use imaging or graphics applications in an effort to understand or represent large amounts of data. Just as frequently, they are frustrated by the long transfer times required for the transfer of these graphics files because of network bandwidth limitations. For this and many other reasons, AMS researchers are in the front ranks of those who are demanding increased network bandwidth. There is also an emergent demand for remote control of experimental equipment, and AMS scientists are working to satisfy this demand (see sidebar). These remote-control facilities will require both high bandwidth and low latency (i.e., almost zero "lag time").

Another critical problem that plagues AMS researchers is that of slow data transfer rates. Typically, ER scientists are geographically remote from their data, and transferring the data they require can take an enormous amount of time. The resulting delays are frequently quite inhibiting and may even render a planned research approach unfeasible. For example, suppose the AMS researchers at LLNL were to run a simulation at the High Performance Computing Research Center at LANL and that that simulation would generate 100 gigabytes of data. Transmitting 24 hours a day at 50 kbps, it would take more than 23 days to transfer the data from LANL to LLNL.

Finally, there is the ongoing and global issue of making the network as transparent and user friendly as possible. At present, the network services linking researchers and the facilities they require do not provide a transparent computing environment. Instead, researchers must often struggle with the network environment itself. To appreciate the effects of this problem, imagine the frustration of the average person if the simple procedure of balancing a checkbook required a thorough understanding of the architecture of the calculator to be used. Just as the solution to that person's real problem (i.e., checkbook balancing) would be delayed by such a requirement, so the solution to the ER scientist's research problem is impeded by the necessity of struggling with the networking environment.

Visual Servoing: A Key “Collaboratory” Technology

The fusion of network communication systems and tool-oriented computing promises to dramatically enhance the collaborative efforts of U.S. scientists. OSC-funded researchers are integrating these two technologies to create what have been called “collaboratories,” or “laboratories without walls.” The ultimate goal of the effort is complete location independence for researchers and research facilities.

As part of this effort, LBL’s Imaging Technologies Group (ITG) is developing a new technology called *visual servoing*, which uses video data as the basis of a closed-loop system for controlling instruments in scientific experiments. In most cases, the instrument so controlled would be the stage or other apparatus on which an object of experimental interest is mounted.

Visual servoing is ideal for applications in which the key control parameters are features (such as shape, location, or color) that are easily captured by video. It is expected to be particularly valuable as an enabling technology for experiments requiring the micromanipulation of samples. In most cases, video monitoring provides the only practical means of collecting real-time positional data on microscopic objects.

Although it is in essence a closed-loop control technology, visual servoing fits into the “collaboratory” scheme because video technology provides convenient mechanisms for remote control. A remote researcher can survey the overall experimental environment by viewing the video images from which the control information is extracted, and on-screen graphics can provide a means of exerting high-level control, enabling the researcher to select an object of interest or a boundary to be tracked.

The drawings on the opposite page illustrate one of the first planned applications of visual servoing. The upper drawing illustrates how visual servoing would be used to make an X-ray beam track the edges of a crystal in a spectromicroscopy experiment. In this implementation, the novel computer vision algorithms developed by the ITG extract the crystal’s boundaries from the raw video data collected by a microscope. The boundary information is then compared with beam-position information derived by other methods. This comparison provides the basis for the control parameters that are fed to the stage control system, which then moves the stage along the path required to keep the beam at the sample’s edge. The lower drawing illustrates how ESnet’s wide-area connections will enable a researcher at a remote site to view the experimental setup and select the initial point for edge tracking.

Other visual servoing applications planned by the ITG include in-situ electron microscopy experiments, video-microscopy-based DNA manipulation, and investigations of the effects of electromagnetic fields on living cells. Besides enabling researchers to carry out novel studies such as these, visual servoing holds tremendous promise for industry, particularly in support of processes such as automated manufacturing and nanofabrication.

The planned implementation of visual servoing on ALS beamline 7.0

Wide-area architecture supporting remote initiation of visual servoing on beamline 7.0

Forecast of Future Requirements

One of the most important future requirements for the AMS program will be to provide the high-speed networking capability required to apply large-scale but geographically dispersed computer systems to the same problem simultaneously. Once developed, this capability would optimize the use of the high-performance computer centers sponsored by the AMS, which are geographically dispersed and house computers of differing architectures. The coordination of such distributed, heterogeneous computer resources offers many advantages for solving problems that are extremely computation intensive. The effectiveness of such distributed interprocessor communication has already been demonstrated in work on complex chemistry problems in the CASA testbed, which operates at over 8000 Mbps.

AMS principal investigators have extended many of the capabilities provided by computer networking, and they continue to push the window in this area. Such "technological push" helps lead the way toward the identification and development of products that will not only serve AMS scientists but will also support the overall ER scientific effort and that of the nation's wider scientific community. Eventually, advanced network products will be disseminated to the general public, in part as a result of formal and informal technology transfer processes that also contribute to the worldwide competitiveness of U.S. networking firms. The work of AMS PIs thus accelerates the development of the National Information Infrastructure (the NII, also known as the Information Superhighway).

The AMS community will be confronted with a wide array of new requirements as the NII evolves. The outlines of some of these new requirements and their probable solutions are quite clear now, while others can only be dimly foreseen. The future challenges for the AMS community are classified below according to how clearly they and their solutions can be delineated at present.

1. *Well-defined new requirements with well-defined solutions that are either site-specific or technology-specific.*
 - Researchers using the Center for Computational Sciences at the Oak Ridge National Laboratory are expected soon to saturate that center's 100-terabit capacity, rendering the existing network bandwidth of 1.5 Mbps obsolete and creating a requirement for at least 45 Mbps. The only alternative to such a network upgrade would be to ship traditional Exabyte tapes around by surface mail, a poor substitute in terms of both time and utility.
 - A typical analysis of data being generated at ORNL for ER researchers at various locations (which include Princeton, La Jolla,

and sites in France and Germany) will utilize the entire capacity of a 1.5 Mbps link for a full ten minutes.

- PNL scientists use the massively parallel IBM computer at ANL. If the output of this computer were visualized via today's ESnet, the researcher would receive the results at the rate of one frame every 20 seconds, or 1/20th of a frame per second. (Compare this to full-motion video at 24 frames per second, which would make it possible for researchers to enjoy full interactivity.)
 - At ANL's Advanced Photon Source, which will come on line within the next 18 months, a typical experiment is expected to generate an 18-megabyte file every 5 seconds. Each data set will contain 1000 files of this size, and a two-day experimental run will produce twenty such data sets. With today's network bandwidths, it would take about 22 days to transfer this amount of data to a remote user.
 - Larger and more widely distributed collaborations will rely heavily on packetized video conferencing.
 - Increased demand for video and audio over ESnet will require improvements in bandwidth and isochronism. Given the bandwidth limitations of today's networks and the delay variations that seem embedded in the network technology, this requirement will challenge the vision, creativity, and fiscal resources of the ESnet community.
 - Using the most sophisticated of the current compression techniques, it will require about 200 kbps bandwidth per user to transmit audio of compact-disk quality. (Note that without such compression techniques, the transmission of such high-quality audio would require about 1500 kbps bandwidth per user. It should also be noted that AMS PIs have significantly contributed to the development of the compression techniques used so heavily in today's computer networks.)
2. *Requirements that are global in scope and certain to become imperative but that are without well-defined solutions.* These are future requirements that are obvious beyond any argument to any informed observer of network evolution. They are also requirements whose fulfillment will entail large-scale alterations in network hardware and operation, often on a global (i.e., planet-wide) scale.
- Increased computing power will allow increased analysis, which will generate ever-growing quantities of data. At the network level, the result will be a demand for vast improvements in data transfer rates, which will require considerable increases in network bandwidth.
 - ER researchers around the world will need greater access to the High Performance Computing Research Centers and the

Supercomputer Access Program Centers. AMS collaborators, in turn, will require network upgrades so that they can provide the facilities required to support this expanded access.

- AMS scientists are increasingly involved in collaborative work, both with each other and with their industrial counterparts. To support the growing number of collaborators, computer networks must become ever more ubiquitous.
3. *New requirements whose impact can be assessed with only moderate certainty at present.* In these cases, we may be able to envisage the outlines of possible solutions without being able to fill in the details with any accuracy.
- As improved network tools make it easier to search the virtual encyclopedia of the world's scientific data, the World Wide Web will become much more useful, resulting in dramatically increased demand.
 - As more and more information is placed on the World Wide Web, and as increasing portions of that information is multimedia data, network upgrades and new types of services will clearly be needed. At present, however, we can only make informed guesses about the scale of those upgrades and the content of new services that will be required.
 - Ever-growing numbers of scientists are involved in ER collaborations worldwide. These scientists will demand increased facilities for creating and delivering multimedia information (including multimedia e-mail).
 - An interactive visual analysis procedure might be repeated many times, as a remote user seeks to discover the optimal virtual location from which to view the data. In the course of this process, the data must be rendered into an image and transmitted to the user iteratively. Using the ESnet links available today, remotely displaying each new image would require one full minute. This iteration rate would exhaust the patience of almost any human user.
 - Researchers will demand the improvement of telework and telepresence tools to serve their needs when they travel to other sites as well as when they are at home but collaborating with other researchers at remote sites. The improvement of these tools will create heavy demands for both network services and increased bandwidth.
 - Researchers will also demand remote control of their experiments. Such remote control systems will require both high bandwidth and low latency. One would expect such systems to provide both full-motion video and precise, high-speed control.

4. *New requirements whose ramifications are almost impossible to predict at present.* These requirements are likely to be driven by “technological push.”
 - As wireless networks evolve, the integration of very small portable and networked devices (of as yet undefined utility) into those networks will have unpredictable effects on network requirements.
 - The implementation of dedicated dark-fiber or OC-48 networks will probably have extensive ramifications that are impossible to specify at present. As networks reach speeds that allow them to be utilized as backplane connections for teraflop-class machines, the scale and pattern of network demand may change considerably, necessitating radical changes in network architecture.
 - The use of networked cyberspace is now conceivable. However, the effects of using 3-D, virtual reality, or other simulation tools over the network are impossible to predict at present.

Notes

1. “The DOE Program in HPCC High-Performance Computing and Communications,” U.S. Department of Energy, Office of Energy Research, Office of Scientific Computing, Washington, DC. March 1993.

Health and Environmental Research

*Robert Weiss
Molecular Biology Institute
University of California, Los Angeles*

*Ray Bair
Environmental Molecular Sciences Laboratory
Pacific Northwest Laboratory*

*John Drake
Mathematical Sciences Section
Oak Ridge National Laboratory*

The Office of Health and Environmental Research (OHER) supports a wide range of projects and programs that span several independent scientific disciplines. The diversity of the HER programs is such that network requirements are best defined project by project and integrated on that basis into the overall assessment of ESnet's user requirements.

The Human Genome Project

Background

Every life form on earth carries the instructions needed for maintaining and propagating itself. These instructions are contained in the nucleic acid molecule DNA or its close relation, RNA and are encoded in four constituent bases, adenine, cytosine, guanine, and thymine. These

bases are associated in varying sequences to form the well-known double helix-structure of DNA. The genetic information, the plan for producing and reproducing life, is contained in the specific sequence in which the bases are associated. A piece of genetic information can therefore be represented as a string of the abbreviations of the four bases, for example, CTAGCCAG.

The genetic information contains specifications for the synthesis of other chemical constituents of life, principally proteins. Proteins are created as chains of their own constituents, amino acids. Each type of protein is thus characterized by a particular sequence of amino acids. However, each protein must also fold into its own unique shape before being able to do its job. Proteins have a wide range of functions, including the chemical synthesis of the nonprotein constituents of life, such as nucleic acids and carbohydrates. Knowing the structure of a protein is enormously helpful in understanding how it performs its function. Currently, this knowledge can be gained only through lengthy and complicated experiments.

The structure of DNA and the significance of its sequence of bases was discovered in the 1950s. During the 1960s, the connection between a DNA sequence and the protein sequence it specified (i.e., the "genetic code") was elucidated, and the first few protein structures were determined. In the 1970s, a practical method for discovering the sequence of a DNA segment was devised. The 1980s brought chemical methods for manipulating DNA and inducing captive bacteria to incorporate specified DNA sequences and produce the related proteins. That decade also saw technological advances in methods of determining protein structures, and these advances led to a rapid increase in the number of known protein structures.

Taken together, these developments created the possibility of studying life in its most intimate details. Consequently, a major reorientation in biological research has focused on the development of tools for structural analysis of proteins and the manipulation of genes. The rise of the biotechnology industry has been based on the use of these tools to affect biological processes.

Through OHER, DOE sponsors research in structural biology and participates in the Human Genome Project, a world-wide cooperative effort to determine the sequence of all the human genetic material. In addition, DOE facilities house several biological databases.

Geographical Description

Thousands of laboratories around the world are engaged in elucidating DNA sequences, and hundreds are involved in determining protein structures. Generally, their results are not only published in the

conventional scientific literature but also stored in network-accessible data banks. The biological research sites and the data bank locations are dispersed throughout the industrialized world. Table 12-1 lists some domestic sites of DOE-funded human genome research.

Table 12-1. DOE-funded Human Genome Research Sites in the U.S.

Houston, TX	Berkeley, CA	Livermore, CA
La Jolla, CA	Tallahassee, FL	Colorado Springs, CO
Los Angeles, CA	Los Alamos, NM	Oak Ridge, TN
Newark, NJ	Providence, RI	Denver, CO
Syracuse, NY	Boston, MA	Rockville, MD
Washington, DC	New York, NY	Upton, NY
Bethesda, MD	Chapel Hill, NC	Salt Lake City, UT
Boston, MA	Ames, IA	Argonne, IL
Tucson, AZ	Ann Arbor, MI	Madison, WI
Richland, WA	Storrs, CT	Tempe, AZ
Palo Alto, CA	Cold Spring Harbor, NY	Baltimore, MD
Philadelphia, PA	Kansas City, KS	Cape Elizabeth, ME
Oak Ridge, TN	Baltimore, MD	Upton, NY
Los Alamos, NM	Bethesda, MD	

In biological data banks, protein structures are stored as lists of three-dimensional atomic coordinates accompanied by annotations indicating the chemical types of the atoms and the experimental and computational procedures that produced the structural data. A database entry for a DNA sequence contains the character string that represents the sequence, along with annotations indicating the species from which it was taken and other information. Figure 12-1 illustrates a typical entry.

TITLE	Transcription from SV40-like monkey DNA sequences				
JOURNAL	Nucleic Acids Res. 12, 4769-4788 (1984)				
STANDARD	full automatic				
BASE COUNT	146 a	169 c	271 g	146 t	
ORIGIN					
1	cccccttcaaa	tctattacaa	ggtgagcgtc	tcgccaaggc	aatgaaatcg caatatgatg
61	tttccattta	ctttggatta	tacgtcatta	taaataattaa	caaataagac tcaaaaagga
121	caccttcggg	taggtcagac	caaagtacaa	aacttgtgtg	tggggctgca gtttgagggc
181	agtgtctgca	gccgtcacat	ggtagcaaaa	cgggtgttaag	cagcgcacga gagtctgcgt
241	cgaccacagc	cagagtccat	gcatcgggag	gttcactcgg	tttgcaaga acgggcaggg
301	catgcacggc	ctgggctcgg	cgggcgggcg	ggcgggccgg	ggcgcagttc cccaggttcg
361	ccactagagg	tcaggaggtg	accgcttcgg	ggctggaaga	cgggcccgtc gtggattggc
421	tagtgccggc	ggagggcggg	gcggagagtg	gggcggggcg	gagagtgggg cggggcgag
481	ttccccaggt	tcgccactag	aggtcaggag	gtgaccgctt	cggggcgggg agactggccc
541	gtcgggggatt	ggctagtgcc	ggcggggggc	ggggcggggg	gcggagggcg ggggtggacgt
601	ggcgccctggt	tgctgacatc	tggaatgact	tttttttggc	atcagatttc ctgtctttgt
661	ggggctgatg	gacccgagta	aagatgcccg	ttcgggggtca	aaggcagagc cgcttctgca
721	gcttctcaaa	gc			

Figure 12-1. Data bank representation of a typical DNA sequence

The LBL Human Genome Center on the World Wide Web

Lawrence Berkeley Laboratory is the site of one of the three primary Human Genome Centers (HGCs) within the DOE Human Genome Program, the others being at LLNL and LANL. The HGCs were established to take advantage of the ability of these national labs to foster interdisciplinary efforts and develop innovative research technologies that would accelerate the realization of program goals. Like the other two HGCs, the LBL HGC is involved in genomic mapping and sequencing. In addition, the LBL HGC has the distinction of being one of two sites nationwide dedicated to production sequencing of human genome DNA.

Another focus of effort at the LBL HGC is *informatics*, the development of information processing systems designed to handle the vast quantities of data generated by the program's mapping and sequencing activities. The Informatics Group provides this computational and software support to external genome research groups as well as to the LBL HGC's own biologists. As one means of addressing this widely dispersed constituency, the Informatics Group makes extensive use of the World Wide Web. For non-HGC scientists and other interested parties, the HGC Web home page provides pointers to brief descriptions of the Center's major projects. In addition, the HGC Web site functions as a server for the database that contains the mapping and sequencing information generated by the Center's researchers. The HGC home page also provides a link to the Biologists' Web Resource page, which displays pointers to a wealth of other biology resources on the Web.

Some HGC scientists and staffers use the Center's Web site as a bulletin-board-like channel for internal communication. The HGC Web site thus provides a place to post notices, the minutes of selected meetings, and other general-distribution items. For the Informatics Group, this Web conduit is a means of communicating bulletin-board style with software developers both within and beyond the LBL group, making it possible to make current bug reports, requests for new features, problem reports, and other software development documents available to all concerned parties.

Members of the Informatics Group have also developed useful extensions to the basic Mosaic interface to the WWW. One such extension made it possible to produce an active outline based on the Web's HTML headers and anchors for the first time. This extension was demonstrated and well received at the Second International WWW Conference in Chicago.

Finally, members of the Informatics Group take an active part in educational outreach activities, using the HGC Web site both as a vehicle for conveying information about the Center and the Human Genome Project and as a means of demonstrating the educational resources available on the World Wide Web.

LBL Human Genome Center

Welcome to the Human Genome Center at Lawrence Berkeley Laboratory

HGC Projects

- Directed Genome Sequencing Project
- Syndb Online LBL Human/Mouse Acedb-style database
- Drosophila Physical Mapping
- Instrumentation Projects
- Informatics Projects
- Human Chromosome Projects

HGC Information

- Introduction to HGC
- Organization
- Publications

Other LBL Servers

- Biologists Web Resources
- LBL Home Page
- LBL Directory
- LBL Library catalog; login: CAT, no password
- Melvyl UC Berkeley Library catalog
- Form-based Melvyl Search page

Other Genomic Information

- DOE Genome Report
- Lookup Biologists at DOE, NIH, NSF, USDA
- NCBI GenBank Server
- National Agricultural Library World Wide Web Server
- Hopkins BioInformatics Web Server
- Quest Protein Database
- Genethon Web Server

for corrections contact HGC Web Master

Information on the LBL Human Genome Center's World Wide Web home page (on-screen formatting differs)

Requirements

The basic requirement is for networks—including ESnet—to provide scientists with prompt and up-to-date access to each other's results. A major point of interest to a researcher who has just determined a DNA sequence is whether a similar or identical sequence has been observed previously. A conventional search of the text literature would be impractical, as the known genetic information amounts to hundreds of megabytes of data. With the rapid and increasing expansion of known genetic information, even database distribution by tape or CD-ROM could not include recent results. Network access allows a researcher to send new results to a database administrator, who can quickly add them to the publicly available store of information, where they can be accessed by other researchers.

Current Network Usage

A variety of methods is used to interact with biological data banks. Depositing newly determined sequences and structures is typically done by e-mail, ftp, or a combination of the two. The most common method of obtaining information from a sequence data bank involves sending an e-mail message containing instructions and a probe sequence to a server at the data bank. After a wait that typically lasts a minute or less, the server replies with an e-mail message containing the search results.

While these canned search procedures suffice for routine use, a laboratory involved in developing novel search methods, for example, might choose to download an entire database periodically. To provide more interactive access, some data banks now support Gopher servers.

The recent explosion of biological information, as well as the discovery of new significance within the information, is driving the development of information storage methods that allow efficient and transparent access, including access to geographically separated databases that are running different database programs. Some data banks are implementing client/server programmatic access to satisfy these requirements.

Both sequence and structural data banks are undergoing exponential growth in information content and access frequency, with both quantities often doubling yearly. Figure 12-2 illustrates the rate at which the Brookhaven Protein Data Bank (PDB), the primary structural database, has been growing in size since the beginning of 1992. Figure 12-3 shows the growth of that facility's ftp traffic over the same period.

Several current trends—improvements in methodology in the biological sciences, increasing interest in the results of biological research, innovations in storage technology and software, and growth in network capacity—create mutually synergistic effects that should sustain such growth in storage and retrieval for the foreseeable future.

Figure 12-2. Growth in size of Brookhaven Protein Data Bank (PDB) from January 1992 through January 1994

Figure 12-3. Brookhaven PDB files downloaded via ftp from January 1992 through December 1993

The Environmental Molecular Sciences Laboratory

The Environmental Molecular Sciences Laboratory (EMSL) is currently being constructed at Pacific Northwest Laboratory under the aegis of the Office of Health and Environmental Research. The EMSL will house 260 permanent and visiting scientists in a 200,000-square-foot facility equipped with world-class instrumentation and computational resources. In addition, the EMSL will be a collaborative research facility, serving both its own internal staff and the wider scientific community at universities, industrial sites, and other government laboratories.

This new facility (see Figure 12-4) will be a key element in ER's expanded response to DOE's growing environmental initiatives. The primary mission of the EMSL will be to develop and refine state-of-the-art methods of investigating molecular processes that are key links in complex environmental processes. To fulfill this mission and to support DOE's broader environmental concerns, the EMSL research program will be based on the following four cornerstones:

- Fundamental and strategic research in the chemical, materials, biological, computational, chemical engineering, and environmental sciences
- Applied research that will focus the results of this fundamental and strategic research onto specific environmental problems
- Research, educational, and training opportunities for the entire U.S. scientific and technical community
- The rapid transfer of knowledge and technology to the commercial sector through early and continuing involvement by industry.

Figure 12-4. An Artist's Depiction of the Future Environmental Molecular Sciences Laboratory

DOE's \$230-million investment in the EMSL includes funds for:

- Constructing the building
- Designing and developing leading-edge instrumentation and software
- Acquiring, installing, and commissioning the research equipment
- Starting up the laboratory
- Supporting the necessary project management functions.

Although the EMSL will not be fully operational until FY 1998, the Office of Energy Research funds a number of research activities that are being carried out at interim EMSL research facilities. These efforts are supported by OHER's Health Effects and Life Sciences programs and by the Chemical Sciences and Geosciences Divisions of OBES.

Requirements

A particular focus of the EMSL program will be the effort to combine the study of fundamental molecular processes (e.g., solvation, reaction, diffusion, absorption, and catalysis in model molecular systems) with simulations of certain complex multispecies, multiphase molecular systems that are crucial in environmental processes. This activity will require a sizable increase in computing and modeling capabilities over the levels achievable with conventional supercomputers and molecular modeling applications. Consideration of the scaling laws for molecular computations indicates that these capabilities must expand beyond current levels by a factor between 100 and 1000. Supporting the use of high-performance computing capabilities at EMSL and elsewhere by widely dispersed collaborators will require appropriate increases in network capacity.

As a collaborative research facility, the EMSL will foster ongoing interactions between researchers at PNL and collaborators at universities, government laboratories, and industrial sites distributed internationally as well as nationwide. These collaborations will depend on ESnet for such activities as geographically dispersed instrument development, multi-institutional software development, network dissemination of research results, and use of the computing facilities at PNL, NERSC, ANL, ORNL, LANL, and Caltech by EMSL researchers at remote sites. As more advanced applications are developed to support collaborative research, remote researchers will be able to control instruments at the EMSL, cooperate in data analysis, and participate in regular discussions about research projects—all without leaving their home sites.

These new collaborative applications will be bandwidth-intensive, often requiring the transport of large volumes of graphical or video data across the network in real time. Applications supporting distributed

modeling and analysis for EMSL projects will also transport much more data than current applications do. It is therefore crucial that the capacity of the certain key network interfaces keep pace with the requirements of the EMSL research program. These interfaces include those between PNL and ESnet, which are now at T1 levels, as well as those between ESnet and the rest of the Internet. An aggressive upgrade policy is needed to ensure that these interfaces support T3 speeds in the near future and still higher speeds as the EMSL becomes fully operational.

Finally, it is important to bear in mind that the EMSL's user base will be widely dispersed, with key collaborators distributed across many non-ESnet sites. Achieving the required end-to-end transport speeds between the EMSL and these collaborators is critical if their research projects are to proceed smoothly and at the desired pace. However, ensuring the required transport speeds may pose a significant challenge, particularly if NSF and other network service providers fail to supply adequate bandwidth.

Global Climate Research

The CHAMMP Climate Modeling Program

The CHAMMP program sponsors research germane to the development of advanced global climate models. The program's acronym, which stands for Computer Hardware, Advanced Mathematics, and Model Physics, reflects the multidisciplinary flavor of the research agenda. One of the primary goals of the program has been to implement the use of massively parallel computers in global climate modeling. The application of this technology would increase both the speed with which climate simulations can be performed and the level of detail that can be included in such simulations. Several state-of-the-art climate models have been implemented in both a data-parallel and a message-passing programming paradigm. A high-resolution ($1/6^\circ \times 1/6^\circ$) ocean model has been developed by LANL in collaboration with the Geophysical Fluid Dynamics Laboratory of the National Oceanographic and Atmospheric Administration. This model has been run successfully in a production mode on a Thinking Machines CM-5 with 1024 processors.

For the first time, these high-resolution simulations are accurately resolving the eddy structures that are present in the oceanic flows and that are important in the poleward transport of heat in the earth's energy budget. One source of such high-resolution simulations will be the Community Climate Model developed by the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. Originally intended to run on Cray's YMP-class machines, the NCAR model has also been successfully implemented on both the Intel Paragon, which has 512 processors, and the Thinking Machines CM-5, which has 1024.

A second focus of the program is the development of coupled climate system models that account for global oceanic and atmospheric circulation in addition to dynamic sea-ice and land-surface processes. The inclusion of all these factors in regionally resolved climate simulations will move the science of climate modeling a step closer to the goal of being able to predict the regional impacts of global climate change.

The Carbon Dioxide Information Analysis Center (CDIAC)

CDIAC provides information and data resources in support of DOE's Global Change Research Program. CDIAC also serves as a repository of global climate change information for a broader international community of researchers, policymakers, managers, educators, and students. The number of requests for CDIAC's data products, information services, and publications has grown over the years. CDIAC's staff now addresses thousands of requests yearly for data and information resources. These requests reflect the multidisciplinary interests of researchers in the physical, life, and social sciences as well as the concerns of requestors from diverse settings in government and business.

Since its founding in 1982, CDIAC has responded to tens of thousands of such requests, supplying the appropriate information resources via surface mail as well as the Internet. During this period CDIAC has also served the wider information-processing needs of the global climate change community. CDIAC is able to serve these wider functions because it is more than a repository of research data. By integrating data from multiple disciplines, CDIAC creates the derived data products, numeric data packages, and computer model packages needed to address complex environmental issues. In addition to evaluating, compiling, and archiving numeric data, CDIAC distributes CO₂-related reports and produces a number of publications, including the newsletter *CDIAC Communications*, which is distributed to more than 9,000 subscribers in 141 countries.

CDIAC also operates one of the World Data Centers for Atmospheric Trace Gases, WDC-A. (The World Data Centers were established by the International Council of Scientific Unions in 1956. The U.S. National Academy of Sciences oversees WDC-A, which includes 12 data centers throughout the United States.) CDIAC's ability to obtain and disseminate global climate change information is enhanced by its role as the operator of WDC-A.

Together with the Martin Marietta Energy System Central Research Library at ORNL, CDIAC also shares responsibility for operating one of the Regional Information Centers established by the International Geosphere-Biosphere Program.

The Atmospheric Radiation Measurement (ARM) Program

ARM is the flagship DOE program in global climate change research. Through careful measurement and detailed modeling of the incoming and outgoing energy fluxes at several critical locations worldwide, the program seeks to reduce the uncertainty associated with modeling clouds and their effect on the radiation budget. ARM's first field site has recently come on-line, and several other field sites are planned (see Figure 12-5).

Figure 12-5. ARM program sites

Networking Requirements for CHAMMP, ARM, and CDIAC

DOE and the CHAMMP program support the nation's global climate research agenda by sponsoring the development of new models and by providing dedicated computer resources for climate simulations that range in temporal scale from decades to centuries. However, these models produce more output than can be stored in a cost-effective manner. For particularly interesting simulations, a high-resolution model run may require 100 Gbytes of archival storage. Local-area high-speed data transfers using high-performance parallel interface (HIPPI) connections between a parallel processor and an archival storage system have been implemented at both LANL and ORNL. Such configurations are typical of the LAN connections required at a computer center to support climate simulations.

The climate modeling programs under the sponsorship of CHAMMP will also require state-of-the-art WAN connections. As the development of coupled oceanic and atmospheric climate models continues, enhanced connections between the national laboratories (in particular, between the High Performance Computing Research Centers at LLNL and ORNL) will become important. A T3 link will support 30–50 researchers using the current generation of high-performance computers at these High Performance Computing Research Centers. The examination of climate model output and the transport of critical files between researchers requires connections with higher speeds than can be achieved over T1 links. The remote display of graphical images could even push T3 links to the limit of their capacity. To allow for the wide-area transfer of large files, major climate model archives will increasingly need to be well connected by high-speed links. These archives include those at the NCAR in Boulder, Geophysical Fluid Dynamics Laboratory at Princeton, and LLNL's Climate Model Intercomparison project. Ensuring access to these centers for the climate modeling community in universities and national laboratories will continue to be important in the future.

ESnet's backbone links are critical in the transport of ARM data. The ARM archive at ORNL is currently receiving data over ESnet's T1 lines. The experimental data are archived at ORNL, as is information related to instrument documentation and data recovery techniques. ARM staffers at that site also disseminate the archived information to the general scientific community. Most of the processing of experimental data occurs at PNL, however. The data are analyzed at PNL, the data products required by the ARM science team are created there, and these data products and other project data are distributed from that site.

The data-transport responsibilities of the ARM centers at ORNL and PNL, like those of CDIAC, make it critical to maintain state-of-the-art network connections between those sites and ESnet. While CDIAC is currently served well by T1 links, the ARM data stream needs speeds greater than T1 links can provide. As the volume of the data from the one currently operational field site increases (as it will when new imaging instruments come on-line, for example) and as new field sites are brought on-line, T1 speeds will become inadequate. Even at the present time, T1 speeds are inadequate for any researcher requesting large volumes of ARM data via the network. The new T3 connections to ORNL provide the network speeds required to support multiple ARM field sites and to handle the increase in data volume expected to be generated by specialized ARM imaging equipment. Over the next one to two years, meeting the increased ARM demand is expected to require network transfer speeds of 150 Mbps.

Human Resources

*Tom Rowlett
Network Management and Engineering Branch
Department of Energy*

The mission of the Department of Energy's Information Resource Management Program is threefold:

1. To provide leadership for all DOE IRM activities, with the goal of ensuring secure, efficient, and effective lifecycle management of the DOE information and information resources that are needed to support the diverse mission of Departmental elements.

2. To advance information technologies that are critical to DOE programs, the national energy strategy, and the national interest.

3. To develop policies, plans, budgets, and standards and to provide services, consultation, oversight, and implementation of appropriate technologies in support of DOE information management activities in a cost-effective manner and in accordance with public law and applicable regulations.

Specifically, the role of HR-4 in the cooperative use of the Energy Sciences network is largely administrative, as HR participates in no major energy research programs. It is the role of the Deputy Assistant Secretary for Information Management (DASIM) to:

- Review and concur (or not concur) as appropriate in non-ER requirements for ESnet use
- Coordinate non-ER requirements with the Office of Scientific Computing for the purpose of budgetary and engineering analysis, documentation, and implementation
- Recover costs, as appropriate, from non-ER DOE elements and reimburse, as appropriate, ESnet operational components
- Coordinate all ESnet-related DOE administrative data communication requirements with OER
- Coordinate operational data communications activities with other ESnet participating members
- Participate in the ESSC and ESCC for the purpose of identifying areas in which ESnet could be enhanced so that it would better satisfy DOE's administrative user requirements
- Participate as appropriate in budget formulation and funding support for DASIM's use of ESnet resources.

Requirements

Transfer of digitized administrative information to DOE sites serviced by ESnet.

Current Network Usage

ESnet provides a pathway to the Internet for non-ER programs so that they can access such systems as:

- Literature-search systems
- Bulletin board systems
- International data banks.

HR has a strong focus on meeting the IRM needs of the Office of the Secretary, which has worldwide communications requirements. ESnet serves as a gateway to the world by:

- Providing the Secretary's Office the means to communicate electronically to energy sites and other governmental agencies worldwide
- Enabling the Secretary's Office to tap into global data banks
- Providing access to DOE for the global community.

Future Requirements

HR's future requirements include:

- Access to the Internet through ESnet

- A system designed to support the budget justification and approval process for construction projects initiated by DOE Field Offices and Management and Operations Contractors.
- A Project Data System (PDS) for the Chief Financial Officer. PDS applications require only file-transfer sessions between users and the central database. Every effort is being made to accommodate the customer's stated objectives for the PDS. Included among these objectives are the use of the TCP/IP protocol suite and the use of the Internet for some connectivity between field sites and headquarters.

In addition, the network must facilitate interactions between the field and headquarters.

Case Study

ESnet services are being used to provide the Secretary of Energy with the capability of exchanging electronic mail with the principal leaders of the various mission areas within the department. This capability includes, for example, providing the directors of the national laboratories with electronic mail access to the Secretary of Energy via ESnet.

Education

"The Information Superhighway" has become a familiar phrase in America. The media has described a glittering array of benefits that it will bring to every home, business, and school in the nation. The President even mentioned it in his State of the Union Address. When Vice-President Gore promotes the National Information Infrastructure, he is most often filmed in a classroom, observing students engaged in an exciting process of discovery by interacting with distant education resources or other students through telecommunications.

Educators are seeking better ways to tap this power for their students. Access to the Internet is being sought by a vanguard of educators at every level. Their vision is that students, teachers, advisors, administrators, and a vast array of on-line educational resources can be linked in a network of opportunity. The media proclaims that the science fiction of the past is today's reality. The concept of long-distance learning has clearly been sold, and expectations are rising.

There is a serious danger that disappointment and frustration will result when these inflated expectations cannot be met in the near term. Technologists realize that the many pilot projects that dramatically show what can be done cannot be scaled up to serve every classroom in the nation without a large expansion of the national networking

infrastructure. There are indications that even the modest number of educational users of such resources as Gopher, WAIS, or World Wide Web are strongly affecting existing networks.

The challenge we now confront is to match the pace of demand by educators and students for on-line services with commensurate increases in the performance of the national networking infrastructure. Well-supported on-line information services are already offered by ESnet. Videoconferencing and audio/video multicast are also technologies that fit well within the vision of the new high-performance learning environments. The use of new instructional technologies will require new skills for teachers, who, in turn, will guide students toward the development of the skills needed to retrieve, organize, and analyze the knowledge that they will be able to access.

ESnet and DOE Actions

With the support of the Energy Research community, ESnet has led the way in introducing emerging technology in a service environment that optimizes usability. ESnet service offerings include a variety of on-line information services. The ESnet technical community is in an excellent position to help accelerate progress toward the level of networking performance that will be required to realize the promise of the Information Superhighway.

Furthermore, DOE has a long-standing tradition of partnership with teachers of math and science throughout the nation. The focus of these programs is to help students by supporting their teachers with tools and training that will enrich curricula and stimulate interest in careers in science and engineering.

The support of teachers, along with their continued training, will be crucial to the success of high-performance learning. As the National Information Infrastructure develops, its new capabilities can be tapped immediately, and the vision of the information superhighway will become a reality. A national community of networked teachers, knowledgeable of and committed to the use of technological innovation, will arise. As teachers become more knowledgeable, and as high-performance learning environments emerge, school boards, community leaders, and parents will be inspired to join in transforming the schoolroom. These tasks will yield results today, while we build the infrastructure required to support the learning environments of tomorrow.

ESnet planning for the deployment of Asynchronous Transfer Mode (ATM) technology represents a cost-effective strategy for achieving the gigabit performance needs of the future.

The DOE National Laboratories conduct a variety of education programs that are directed toward the full spectrum of education levels, from kindergarten to the university level. The following list is not comprehensive, but it indicates the kinds of programs being pursued at the laboratories.

- Lawrence Berkeley Laboratory's Microworlds, Hands On Universe, and Whole Frog projects
- Work by the Princeton Plasma Physics Laboratory to assist the state of New Jersey in providing Internet access to K-16 schools
- Stanford Linear Accelerator Center's work with the Public Broadcasting System's Learning Link
- The New Mexico school administrators' workshop on technology run by Sandia
- Los Alamos National Laboratory's Supercomputing Challenge Program
- Ames Laboratory's program to make a 128-node parallel processor available to K-12 schools
- Oak Ridge National Laboratory's education network
- Fermilab's Saturday Morning Physics Talks, which reach some 9,000 teachers and 40,000 students over telecommunications networks
- Idaho National Engineering Laboratory's water-quality project
- Pacific Northwest Laboratory's ESnet server for local schools
- Lawrence Livermore Laboratory's National Energy Research Supercomputer Center, which provides dedicated Cray Supercomputer resources to over 100 classes and 3500 users annually
- The collaboration of Oak Ridge, Ames, and Sandia National Labs, along with K-12 schools and universities, in running the Adventures in Supercomputing program.

Toward the Virtual Workplace: New Modes of Scientific Collaboration

*Ray Bair
Environmental Molecular Sciences Laboratory
Pacific Northwest Laboratory*

*Stu Loken
Information and Computing Sciences Division
Lawrence Berkeley Laboratory*

In this era of explosive growth in technical and scientific information, optimizing scientific productivity and innovation depends on ensuring efficient access to a wide range of knowledge resources that are generated by colleagues in institutions distributed worldwide. As scientists address increasingly complex problems, they often develop one-of-a-kind instrumentation or computer models, and it is often crucial for colleagues to be able to leverage the special expertise and unique resources produced as byproducts of this developmental process.

The growing requirements of scientific collaborations and research facilities thus drive the creation of new models for collaboration. In order to collaborate, scientists must communicate effectively and be able to share many aspects of their work. Important scientific communication may range from the kind of casual discussions typically held in hallways to exchanges of large volumes of experimental or computational data.

The facilities used today by DOE scientists include leading-edge laboratory instruments, sophisticated instruments at large particle accelerators, large-scale computing facilities, and a host of standard analytical and computing tools. Shared research facilities have long had a major role in DOE research, and with the growing importance of multi-institutional projects, site-to-site communication becomes increasingly important. In fact, the success of many DOE projects depends on the rapid, dependable, and easily established communications that computer networks have made an essential part of daily scientific work.

The character of scientific collaboration also changes as its day-to-day importance grows. DOE collaborations increasingly involve not just national efforts but international groupings as well. Such collaborative groups may grow beyond the traditional scientific research community to include active participants from industry. Like the traditional scientific community, this extended collaborative family must communicate well to work effectively.

Steady advances in electronic communications are producing important new methods of communication. In the early days of networking, electronic sharing meant that scientists used electronic mail to send and receive text documents. Over the last decade, we have progressed beyond electronic mail to faster forms of communication with greatly increased capabilities. But the difference between e-mail exchanges and our present communications tools is dwarfed by the gulf that separates our current networking capabilities from the communications technologies now on the horizon. These emerging technologies will result in new opportunities to enhance the communication of scientific information, to improve the productivity of scientific collaboration, and to increase the effectiveness of technology transfer efforts.

All of these changes are creating a world in which location is largely irrelevant. Very soon scientists will be able to work with other scientists, obtain information, or utilize research facilities regardless of where those resources are physically located. The emergence of this "virtual workplace" will alter and enhance the capabilities of the scientific process as a whole.

Enabling Technologies

A number of key technologies are now coming together to make the virtual workplace a reality. These enabling technologies include:

- A variety of facilities for high-speed computation involving large data files
- Mass-storage facilities linked to these computational resources by the network
- Distributed computing frameworks that make it possible for numerous applications running on various platforms to share data
- High-speed wide-area networks
- Videoconferencing and electronic whiteboard applications that create a "telepresence" that approximates face-to-face interaction for their users.

These technologies, together with other critical enabling factors, such as the development of widely accepted network and windowing standards and the advent of low-cost desktop and portable computers, are making new forms of collaboration available worldwide.

Shared Computing and Mass Storage Facilities

Although the computer systems used by individual scientists are more powerful than ever, shared computing and mass storage facilities continue to play an important role in DOE research. Included in this category are general-purpose supercomputing facilities such as the NERSC at LLNL and the High Performance Computer Research Centers at LANL and ORNL as well as specialized facilities at other DOE laboratories. Such high-performance computing facilities have their own high-capacity mass storage facilities on-site or are linked to such storage facilities via the network. As our computing capabilities grow, so will our need for fast, scalable mass storage that can handle the voluminous amounts of data and information generated by advanced applications. Because much of this data will be multimedia data and because much of it will need to be accessed in real time, appropriate service guarantees and flow control will be required.

The network is often the exclusive avenue for access to high-performance computational resources such as those mentioned above. Emerging capabilities for exploiting these resources include the ability to select a computer architecture that is well suited to a particular aspect of a research problem, the ability to access that computer from any desktop, and the ability to capture and analyze the results of computations as they proceed.

Scientists can also create computing environments that combine computation and information resources from a number of local and remote shared computing facilities into a virtual supercomputer. When the scientists in question no longer need this virtual computer, the allocated facilities are released and again become available to others. A number of advances in distributed computing and networking have made it much easier to build and manage these virtual systems. Software tools like the ORNL-developed PVM (Parallel Virtual Machine) and P4 (Portable Programs for Parallel Processors) from ANL provide frameworks for distributing a single computation across a set of heterogeneous machines.

Advances in high-speed networks are also key enabling factors in high-performance computing. Distributing large computations among several supercomputers can bring to bear greater computing power than is available from any single computer. This approach also makes it possible to use the most suitable computer for each step of the task. However, today's fastest wide-area networks typically operate at only 45 Mbps, and communication at this rate creates a significant bottleneck when the network is linking computers that generate results at a hundred times that rate. The solution to this problem lies in the development of gigabit networks (i.e., networks capable of carrying billions of bits per second).

The CASA Gigabit Network Testbed, a collaboration involving many research institutions and industrial organizations, has the goal of demonstrating leading-edge distributed applications that run on multiple high-performance computer systems linked by a gigabit network. The CASA wide-area gigabit network links LANL, the San Diego Supercomputer Center, the California Institute of Technology, and Caltech's Jet Propulsion Laboratory via high-speed commercial communications lines and first-of-a-kind computer hardware and software. The CASA project developed new parallel programming techniques to support the construction of scientific applications that utilize multiple parallel computers of varying design. In addition, the HIPPI-SONET gateway developed at LANL provided an elegant solution to the problem of ensuring compatibility between the Synchronous Optical Network (SONET) standards that have been adopted by long-haul telecommunications carriers and the local networks that typically serve high-performance computer facilities.

One application employing all of these capabilities is a quantum chemical reaction dynamics model. This model partitions its work between supercomputers at Caltech and LANL by assigning different stages of a computation to each system and transferring intermediate results over the gigabit connection. The aggregate supercomputing power enables researchers to make calculations over a wider range of reaction conditions than was previously possible.

The use of these virtual systems is not restricted to modeling and simulation. Data collection computers that collect experimental data from high-energy accelerators, light sources, or high-field magnetic resonance instruments can be integrated with computers at remote locations that perform data analysis, visualization, and storage. This integration not only allows scientists to share data, it also enables them to make better use of precious instrument time. While an experiment is running, they can use shared high-powered computers to analyze incoming experimental data, look at preliminary results, and then modify experimental parameters to collect better data.

Advanced Research Databases

The same advances in electronics that made it possible to model more complex physical and chemical systems have also made it possible to develop new generations of advanced research instrumentation. Computer controlled instrumentation and digital data acquisition systems have become integral elements in the apparatus required to perform advanced research. In many cases, these instruments generate large volumes of data. Although much of the data remains with individual scientific groups for analysis and interpretation, a significant fraction is now “published” in databases that comprise the authoritative records in specific disciplines. These databases also serve as a collective repository for data produced by numerous researchers in the same field. For example, data from thousands of researchers working on the Human Genome Project are kept in large databases that are accessible from anywhere in the world. Similarly, the Review of Particle Properties database contains information on all of the elementary particles and their interactions. Networks are the medium through which these national databases are constructed and used.

The growing impact of databases, however, creates a number of challenges for the scientific community. Foremost among them, perhaps, is that of providing researchers a uniform means of accessing a diverse set of data that are maintained by a number of different groups. Ensuring such access is crucial, as researchers cannot be experts on all the collections relevant to their fields of work, and a given researcher may not even know what is available and when. Emerging collaborative research tools are just beginning to provide the distributed computing environments that are solving the problem of knitting together individual research communities and their data.

Distributed Computing Frameworks

As the need for distributed computing and data sharing grows, so does the need for increasingly general frameworks for software systems. Traditionally, each software application, research database, visualization

tool, or data management tool had its own unique set of input parameters and data file formats. This heterogeneity meant that a lot of “glue” in the form of custom software was required to allow programs to share data. Tools such as PVM provide communication mechanisms for one important class of distributed systems, but researchers now need much more general software frameworks. These new frameworks must be able to conjoin a variety of applications and their data while managing the execution of complex, multistep computational processes.

The architecture of a computing framework is the key to its success, and advanced software engineering techniques are producing promising new approaches. Among the factors that determine the success of such a framework are its ability to manage data consistently, its ability to add new data items easily, and its ability to use existing code to develop derivative programs. It is just these properties that are the strengths of object programming. Hence, object design, object-oriented data models, and object databases are playing a growing role in the construction of advanced computing frameworks.

One key to the development of such frameworks is the creation of detailed specifications for standard object-oriented software interfaces. For the scientific world, the standardization of such interfaces will mean that many of the features of object-oriented programming can be extended to distributed application environments. Among other benefits, such object-oriented programs will provide a high degree of transparency for work that utilizes computational and data resources across the network. Objects can be reused in new applications, thus encouraging common solutions to problems of data handling and user interfaces. Emerging standards for distributed objects (standards such as the Common Object Request Broker Architecture, or CORBA) will make it easier to build distributed-object systems on a wide variety of platforms.

Telepresence Tools: Workstation Videoconferencing and the Electronic Whiteboard

Workstation-based videoconferencing (see Figure 15-1) will be a critical element in future collaborative work environments. Using workstation video tools, scientists can carry on face-to-face discussions much as they would if they were in the same office or laboratory and attend meetings without traveling. Workstation videoconferencing will also allow experimenters to observe the operation of remote research apparatus and detect details that would not be apparent from the data stream alone.

Figure 15-1. Workstation Videoconferencing

Developing videoconferencing tools for use on packet networks like ESnet has been a major challenge for network researchers, and DOE-funded computer scientists have had an important role in developing this capability. To successfully transport video information, it is necessary to convert the video input into network packets, transmit that stream of packets in a rapid, orderly manner, and route the packets to the proper destinations. Finally, the videoconferencing application must correctly reconstruct the video image at the receiving workstation.

The first major problem for researchers to solve was that of matching data flow to available network bandwidth. Video cameras produce megabytes per second of raw data. Fortunately, the amount of unique information in a given frame of video data is considerably smaller. Data compression techniques therefore make it possible to transmit images of reasonable quality at a small fraction of the data rate produced by a video camera. Nevertheless, sending individual copies of video packets to all recipients of a video conference would rapidly swamp network links.

To address these bandwidth concerns and to create an environment that can be scaled to very large collaborations, researchers developed

conferencing tools that employ multicast protocols. Multicast technology ensures that only one copy of a video conference is transmitted over any particular physical network link. This means that a video conference with hundreds of participants is possible on existing networks. However, network bandwidth still limits the number of different video conferences that can simultaneously traverse the same links.

Beyond the problem of bandwidth, the greatest challenge for researchers was dealing with the variable-delay problem inherent in network packet transmission. Because video data is real-time information, delays and transmission problems frequently resulted in poor picture quality and lost video frames. Fragments of a conversation would arrive at random times, and sometimes in random order. To address this problem, time stamps were added to packets to ensure that they would be reassembled in the correct sequence and with the correct timing.

The videoconferencing tools now in use on ESnet were designed to scale to arbitrarily large meetings. A session-management tool provides users with a list of meetings and connects them to a selected meeting without interrupting the ongoing discussion. These videoconferencing tools also provide interoperability among most of the UNIX-based workstations used in the scientific community. However, the current workstation video applications do not extend this service to Macintoshes or PCs. (Since these systems do not yet support preemptive multitasking, they cannot interoperate with the UNIX systems.) It is expected that continued DOE-funded development of packet-based videoconferencing and the underlying protocol infrastructures will make significant contributions to the development of this technology and speed commercial development of reliable teleconferencing applications.

Another emerging telepresence tool is the electronic whiteboard, a shared drawing and writing surface that permits multiple collaborators to write or type on the same surface, just as they would if they were in the same room. The electronic version offers a number of advantages over the physical whiteboard; drawings or documents from a computer can be displayed on it, more than one person can draw on it at once, and it can keep track of each participant's entries separately.

New Opportunities for Traditional ESnet Customers

Within the Energy Research community, the initial users of new computing and information technologies will be ESnet's traditional customers, namely, scientists at the DOE laboratories and university-based scientists and their colleagues. Many of these researchers work permanently at DOE research facilities, while others do so periodically as a function of their participation in specific scientific programs. ESnet provides all of these scientists with shared access to computational

resources and data from ongoing experiments. The network also supports research collaborations involved in the design of new experiments and instruments. For scientists involved in these design efforts, ESnet makes it possible to control and monitor critical pieces of apparatus, so that problems with experimental equipment can be corrected very quickly.

The next generation of experiments will utilize new computing and communication technologies that will permit better feedback between instrument design, modeling, and operation. Systems under development will support more data sharing as well as remote control of more types of apparatus. It will even be possible to fully control some types of scientific experiments from remote locations across the country or around the world. Imaging technologies such as computer vision will assist in experimental control and feature extraction by recognizing the specific parts of an image that are of interest in the experiment.

The experimental process will be further enhanced by closer coupling between theory and experiment, as distributed modeling systems and distributed instrument-control systems are integrated, allowing scientists to validate assumptions and optimize data collection in real time. Likewise, the next generation of modeling and simulation software will make increasing use of networks to support geographically distributed software development, the sharing of model data, and wider use of special-purpose computing facilities.

New Channels for Scientific Communication

It is important to bear in mind that the primary product of DOE research is information. The traditional formats for information distribution in science have been the scientific paper and the technical report, and their medium was the printed page. However, these traditional means of distributing information impose limits on the types of information that can be disseminated, on the audience that has access to that information, and on the ways in which the information can be retrieved and applied. The implementation of electronic information servers that use the network for distribution transcend these problems. Such servers will serve as a principal medium for conveying the value of DOE's national missions to a much broader audience.

Multimedia-capable servers can present DOE information in many different forms, providing greater insight into the scientific research being described. Such multimedia technologies, which combine color images, video animation, and audio, are powerful tools for better presenting information related to DOE's scientific missions.

The delivery of multimedia requires a ubiquitous network. For most DOE scientific and technical data, CD-ROM and other single-user storage

technologies are too low in capacity and too slow to deliver the quality and breadth of information required by DOE. The World Wide Web (WWW), originally developed by the High Energy Physics community in Europe, is currently performing this function on the Internet. The Web's worldwide network of information servers is today's technology of choice for delivery of multimedia information. For millions of users, the Mosaic graphical interface provides a means of accessing and browsing on the Web.

The WWW is the forerunner of even more powerful and versatile information infrastructures that are expected to revolutionize commerce and entertainment as well as the delivery of scientific information. Such tools will allow scientific information to be distributed even more effectively to scientists, engineers, stakeholders, and clients.

Today the communications revolution spawned by the development of high-speed networking is beginning to affect the lives of a much wider community. Millions of employees in the United States already use telecommunications as an alternative to travel. Even the sporadic use of telecommuting can have an impact on energy utilization as well as on employee productivity. More intensive use of telecommunications can lead to the development of new services and new economic opportunities. In the scientific domain, network transmission of professional meetings, workshops and other scientific discussions are already moderating travel expenses and broadening the dissemination of research results.

As the scientific community continues to work more closely with industry, telecommunications can make unique facilities available to our national industrial base. For example, electron microscopes and synchrotron radiation facilities can be used to study impurities in samples, and biological databases can aid the search for new pharmaceuticals.

Another growing constituency is the educational community, which has been an important focus of DOE's mission for many years. The same tools being developed for scientific collaborations also offer a major opportunity to fundamentally change the way that DOE scientists interact with students and teachers at all levels. Today, projects such as the Hands On Universe are permitting students to experience the excitement of scientific research by allowing them access to data from on-line experiments. In the future, videoconferencing will provide students with "virtual field trips" in which they will meet with leading scientists who work on a variety of challenging problems. This kind of experience—coupled with access to experiments, computer systems, databases, and document servers—will not only improve the quality of education but also encourage the choice of science as a vocation.

Conclusion

As the network has become the web connecting the worldwide scientific community, much of our computer hardware and experimental apparatus has been transformed into a collection of peripherals on the network. The fusion of new computing and communication technologies will enable us to transcend our current methods of collaboration and create a new work paradigm in which the laboratory, the computer center, and the library become as global as science itself.

ESnet Tomorrow

Today ESnet provides a wide spectrum of services to its user base. The dramatic growth in usage of existing as well as emerging services underlines the need for new services and significant improvements in bandwidth and server capability. ESnet's goal, as always, is to measurably enhance the DOE research environment through these improvements.

In pursuit of this goal, the Energy Sciences Network is about to begin a transition from a dedicated network consisting of leased lines connecting DOE laboratories and associated universities to commercially provided network transport services. These services will provide a matrix of "virtual" connections between all ESnet sites. The new technologies associated with these network services define a radically new approach to the provision of wide area networking. Eventually these technologies will find their way into the local-area network environment and to every desktop computer system.

Because ESnet is still in the process of procuring these advanced networking services, this Program Plan cannot provide any technical or financial details relating to ESnet's plans to utilize these services. One can assume that the ESnet deployment will eventually go beyond T-3 speeds and into the gigabit range, if necessary. It is also assumed that the ESnet

implementation may support advanced applications that can take advantage of the isochronous capabilities of the new networking technologies.

ESnet has engineered its topology with more diversity than any other mission oriented network. Its independent connections to virtually every regional network in the country will minimize disruptions as the rest of the Internet adjusts to the commercial replacement of the NSFnet. Additionally, ESnet intends to attach to NSF's Network Access Points as needed to maintain connectivity to the research and education communities on the Internet.

For now and into the foreseeable future, ESnet's user-driven ethos is prepared to meet the growing demands created by the next generation of scientific applications and deploy new technology as needed to meet the requirements of DOE's research programs.

Appendix

Acronyms and Abbreviations

AFS	Andrew File System
AMS	Applied Mathematical Sciences
ANL	Argonne National Laboratory
ARPA	Advanced Research Projects Agency
AS	Autonomous System
BITnet	"Because It's Time" Network
BNL	Brookhaven National Laboratory
CADD	Computer Aided Design and Drafting
CCIRDA	Coordinating Committee for Informatics Research, Development, and Application
CCIRN	Coordinating Committee for International Research Networking
CDIAC	Carbon Dioxide Information Analysis Center
CEBAF	Continuous Electron Beam Accelerator Facility
CERN	European Organization for Nuclear Research
CHAMMP	Computer Hardware, Advanced Mathematics, and Model Physics
DARPA	Defense Advanced Research Projects Agency
DCCC	Distributed Computing Coordinating Committee
DNS	Domain Name Service
DESY	German Electron Synchrotron Laboratory
DOD-IP	Department of Defense Internet Protocol

DVC	Desktop Video Conferencing
EDA	ITER Engineering Design Activity
EM	Environmental Waste and Restoration Management
EMSL	Environmental Molecular Sciences Laboratory
ER	Energy Research
ERSUG	Energy Research Supercomputer Users Group
ESCC	ESnet Site Coordinating Committee
ESSC	ESnet Steering Committee
EXERSUG	Executive Committee, Energy Research Supercomputer Users Group
FDDI	Fiber Distributed Data Interface
FE	Fusion Energy
FEPG	Federal Engineering and Planning Group
FIX-E	Federal Interagency eXchange-East
FIX-W	Federal Interagency eXchange-West
FNAL	Fermi National Accelerator Laboratory
FNC	Federal Network Council
FTP	File Transfer Protocol
GA	General Atomics, La Jolla, CA
GIX	Global Internet eXchange
GRM	Grumann Aerospace, Princeton, NJ
HEP	High Energy Physics
HEPAP	High Energy Physics Advisory Panel
HEPnet	High Energy Physics Network
HFBR	High Flux Beam Reactor, BNL
HFIR	High Flux Isotope Reactor
HIPPI	High-Performance Parallel Interface
HPCC	High Performance Computing and Communications program
HPCCIT	High Performance Computing and Communications Information Technology
HR	Human Resources and Administration
HRC	HEPnet Review Committee
IEPG	International Engineering and Planning Group
IETF	Internet Engineering Task Force
IHEP	Institute for High Energy Physics, Protvino, Russia
IITA	Information Infrastructure Technology and Applications
IITF	Information Infrastructure Task Force
INEL	Idaho National Engineering Laboratory
INFN	Italian National Institute for Nuclear Physics Network
IP	Internet Protocol
IPMS	ITER Integrated Process Management System
IPng	Next-Generation Internet Protocol
IPNS	Intense Pulsed Neutron Source, ANL
IRM	Information Resource Management
ISDN	Integrated Services Digital Network
ITER	International Thermonuclear Experimental Reactor Project, La Jolla, CA
JCT	Joint Central Team sites, San Diego, Naka, and Garching
JPL	Jet Propulsion Laboratory

KEK	Japanese National Laboratory for High Energy Physics
LAN	Local-Area Network
LANL	Los Alamos National Laboratory, NM
LBL	Lawrence Berkeley Laboratory, CA
LHC	Large Hadron Collider, CERN
LLNL	Lawrence Livermore National Laboratory, CA
M&O	Management and Operating Contractors
Mbone	Multicast backbone
MCU	Multiway Conference Unit
MFEEnet	Magnetic Fusion Energy Network
MIPS	Million Instructions per second
MIT	Massachusetts Institute of Technology, Boston, MA
MOU	Memorandum Of Understanding
NAS	National Academy of Sciences
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NERSC	National Energy Research Supercomputer Center
NFS	Network File System
NIC	Network Information Center
NII	National Information Infrastructure
NMFECC	National Magnetic Fusion Energy Computer Center
NOC	ESnet Network Operations Center, LLNL
NSF	National Science Foundation
NSI	NASA Science Internet
NSLS	National Synchrotron Light Source, BNL
NSTC	National Science & Technology Council (formerly FCCSET).
OBES	Office of Basic Energy Sciences
OER	Office of Energy Research
OHER	Office of Health and Environmental Research
ORNL	Oak Ridge National Laboratory
OSC	Office of Scientific Computing
OSI CLNP	Open Systems Interconnection Connectionless Network-layer Protocol
OSI	Open Systems Interconnection
PDB	Protein Data Bank
PI	Principal Investigator
PNL	Pacific Northwest Laboratory, Richland, WA
PPPL	Princeton Plasma Physics Lab, Princeton, NJ
PVM	Parallel Virtual Machine
RHIC	Relativistic Heavy Ion Collider
SCIE	Scientific Computing Information Exchange
SCS	Scientific Computing Staff
SLAC	Stanford Linear Accelerator Center, Palo Alto, CA
SNL	Sandia National Laboratories
SNLA	Sandia National Laboratory, Albuquerque, NM
SNLL	Sandia National Laboratory, Livermore, CA
SNO	Sudbury Neutrino Observatory

ESnet Program Plan 1994

SONET	Synchronous Optical Network
SSC	Superconducting Super Collider, Waxahachie, TX
SSRL	Stanford Synchrotron Radiation Laboratory
TCP	Transmission Control Protocol
TFTR	Tokamak Fusion Test Reactor
TPX	Tokamak Physics Experiment
UTA	University of Texas at Austin
VCS	Video Conferencing Service
VCSS	Video Conferencing Service Scheduler
WAIS	Wide Area Information Server
WAN	Wide-Area Network
WDC	World Data Centers
WWW	World Wide Web