ESnet Program Plan 1998

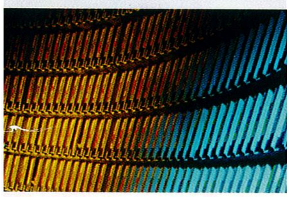

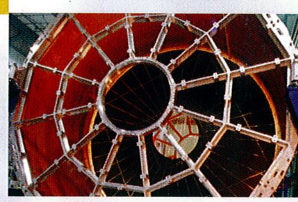

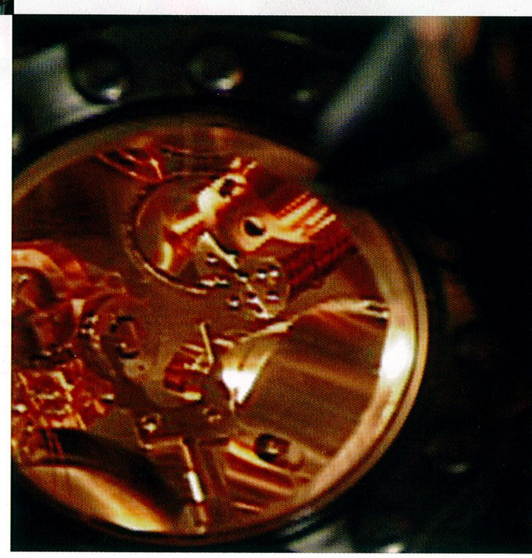

DOE/ER-0719

ESnet Program Plan

March 1998

U.S. Department of Energy
Office of Energy Research
Mathematical, Information,
and Computational Sciences Division
Germantown, MD 20874

DOE/ER-0719

March 10, 1998

Martha Krebs, Director, Office of Energy Research
David Nelson, Associate Director for OCTR
Daniel Hitchcock, Acting Director, MICS Division
George Seweryniak, ESnet Program Manager

I am pleased to submit to you the 1998 Energy Sciences Network (ESnet) Program Plan, on behalf of the ESnet Steering Committee. The Steering Committee is made up of representatives of Energy Research programs and other DOE programs that use ESnet. It is charged with identifying and prioritizing network requirements and with reviewing implementation plans and resource allocation. The close cooperation and communication between ESnet providers and users that results from the Steering Committee process is one of the strong "secrets" of the success of ESnet in providing networking that is reliable, innovative, and above all efficiently targeted at the mission requirements of its user programs.

This document was requested by the Mathematical, Information, and Computational Sciences Division of DOE's Office of Energy Research. It is itself the result of a broad collaborative process with contributions from many principal investigators across Energy Research and the Department of Energy. Their willingness to devote their time to producing this Program Plan has been informed by their understanding that networking is absolutely critical to the success of their research and that their reliance on the network is growing rapidly.

The Program Plan provides a broad summary of the research programs served by ESnet and uses that to outline each program's present and future networking needs. The Plan also describes the organization and processes that ESnet uses to identify and meet these needs. I hope you will find in this document a vivid reminder both of the many ways in which networking provides "glue" to DOE's widely distributed programs and of the foresight with which ER and DOE have provided this particularly enabling infrastructure to its researchers. The Plan makes it quite clear that today's network will not be adequate tomorrow and that meeting the dynamic needs of DOE research will continue to be a challenge.

The ESnet Steering Committee appreciates the support and priority given to ESnet in the past and looks forward to a continued partnership in this critical component of DOE programs. Please let us know of any specific amplification of this document or other information that would assist you in supporting this crucial network.

Sincerely,

Lawrence E Price

Lawrence E. Price, Chairman
ESnet Steering Committee

The Energy Sciences Network Steering Committee (ESSC) 1998

Raymond A. Bair
Deputy Director, EMSL
Pacific Northwest Laboratory
Battelle Boulevard, MS K8-91
Richland, WA 99352
RayBair@pnl.gov

Greg Chartrand
Pacific Northwest National Laboratory
Battelle Boulevard
Richland, WA 99352
greg@pnl.gov

Martin Greenwald
Principal Research Scientist
Massachusetts Institute of Technology
Plasma Fusion Center
175 Albany Street
Cambridge, MA 02139
g@psfc.mit.edu

Dale Koelling
Materials Science Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439
koelling@anl.gov

Jean-Noel Leboeuf
Group Leader
Fusion Energy Division
Oak Ridge National Laboratory
Bldg. 9201-2, MS 8071,
P.O. Box 2009
Oak Ridge, TN 37831-8071
leboeuf@fed.ornl.gov

Jim Leighton
Head, Networking and
Telecommunications Department
Lawrence Berkeley National Laboratory
One Cyclotron Road, MS 50B-4230
Berkeley, CA 94720
jfl@es.net

Sandy Merola
Deputy Head
Computing Sciences Directorate
Lawrence Berkeley National Laboratory
One Cyclotron Road, MS 50B-4232
Berkeley, CA 94720
merola@lbl.gov

John F. Morrison
Los Alamos National Laboratory
Mail Stop B260
P.O. Box 1663
Los Alamos, New Mexico 87545
jfm@lanl.gov

Larry Price, ESSC Chair
Director
High Energy Physics Divison
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 362
Argonne, IL 60439
lprice@anl.gov

Thomas M. Rowlett
Office of Information Management/
Engineering Services Team
Department of Energy
HR-423
Washington, DC 20585
thomas.rowlett@hq.doe.gov

Rick L. Stevens
Mathematics and Computer Sciences
Argonne National Laboratory
9700 South Cass Avenue
MCS 221, A-221
Argonne, IL 60439
stevens@mcs.anl.gov

Walter Toki
Professor of Physics
Dept. of Physics
Colorado State University
AR27 Engineering Bldg.
Ft. Collins, CO 80523
toki@lamar.colostate.edu

Jimmy W. Voyles
Manager, ARM Data & Science
Integration
Pacific Northwest Laboratory
Battelle Boulevard
Richland, Washington 99352
jimmy.voyles@arm.gov

Robert Weiss
Programmer/Analyst
MBI 205
University of California, Los Angeles
Los Angeles, CA 90095
weiss@ucla.edu

Roy Whitney
Head of Computing
Thomas Jefferson National
Accelerator Facility
12000 Jefferson Avenue
Newport News, VA 23606
whitney@jlab.org

Chris Witzig
Physics Department
Brookhaven National Laboratory
Building 510-C
Upton, New York 11973
witzig@bnl.gov

Contents

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 known today, and ESnet continues to contribute to the future of networking through its utilization of leading-edge technologies and its participation in national and international research and development projects. These projects bring ESnet and the entire Internet to the threshold of major changes in technology, operations, and administration.

ESnet is a service-oriented network chartered by the Department of Energy (DOE) to support forefront 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 enhance the research capabilities of 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. DOE programs represented on the ESSC are Basic Energy Sciences; Biological and Environmental Research; Defense Programs; Fusion Energy; High Energy Physics; Human Resources; International Nuclear Safety; Mathematical, Information, and Computational Sciences; and Nuclear Physics. 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 the advancement and sharing of 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 national multi-agency initiatives that apply technology to the opportunities of today. The majority of this Program Plan, in fact, is devoted to descriptions of the requirements of the DOE programs that rely on ESnet. This Program Plan is the product of the members of the ESnet Steering Committee and other principal investigators who represent the broad user base of ESnet.

At its inception, ESnet consisted largely of leased lines and hardware that connected ER sites together in a complex network web. Network evolution has seen these dedicated communications facilities replaced with a publicly shared communications infrastructure provided by private telecommunications vendors. ESnet has 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 from ER 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.

In this era of unprecedented political and social interest in computer networks, ESnet faces many challenges and opportunities. By capitalizing on its strengths—outstanding service and reliability, leading-edge technologies, and collaboration with users—ESnet will maintain its leadership in the field of networking and its increasingly important role as an essential tool for scientific research.

Current information about ESnet is available on the World Wide Web at http://www.es.net/.

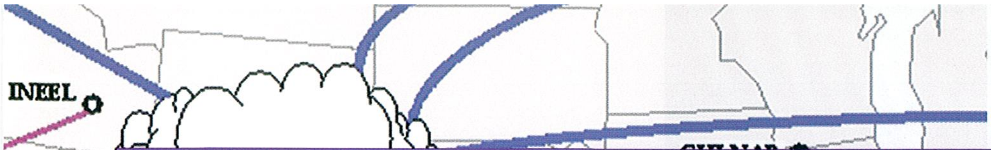

ESnet Charter and Goals

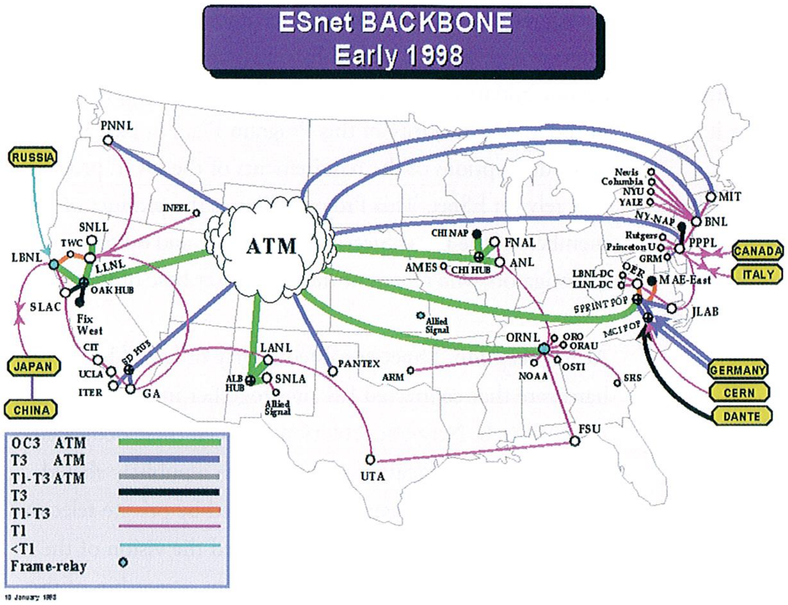

The mission of the Energy Sciences Network (ESnet) Program is to provide highly capable and reliable communications infrastructure and leading-edge network services that support DOE's missions. The program emphasizes advanced network and distributed computing capabilities needed for forefront scientific research and other DOE programs, thus enhancing national competitiveness and accelerating development of future generations of communication and computing technologies. Participating DOE programs and their partners actively contribute to program planning and implementation, and collaborate on development of new technologies that will drive upcoming generations of research capabilities.

As a mission-oriented program, ESnet is organized to provide the best possible networking for DOE programs. Many of the participating DOE programs rely fundamentally on the capabilities of the network to enable their research functions. For them, advances in network capabilities translate directly into advances in research capabilities: rate of progress, access to instruments, better insights, more efficient use of time for researchers and for their instruments, all of which can be summarized as research power. This vital connection has led ESnet to pursue a number of goals whose implementation is described in this Program Plan:

- a reliable production network
- close coupling to programmatic requirements
- continued improvements in network service and related applications targeted at the emerging needs of the programs
- active interaction and coordination with ESnet sites to optimize service and performance
- effective interagency and international coordination and cooperation.

Why ESnet Is Vital to DOE

As an agency at the core of fostering and funding research in the United States, the Department of Energy finds it imperative to implement technology that will:

- enhance the rapidity at which the research culminates
- broaden the community that participates in the research
- make efficient use of the experimental facilities
- stretch the funding to maximum benefit.

Computer networking is a proven technology that accomplishes these goals, and ESnet is paramount in the world as a network that supports researchers. DOE is therefore extremely fortunate to have this essential ingredient for successfully achieving the Departmental mission in the next decade.

ESnet History and Accomplishments

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 FE community had been served by the Magnetic Fusion Energy Network (MFEnet), which was launched in 1976 as a result of the 1974 opening of a dedicated FE supercomputer center at Lawrence Livermore National Laboratory (LLNL). In order to make use of the new National Magnetic Fusion Energy Computer Center (NMFECC), 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 NMFECC 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 National Laboratory (LBNL). In the early 1980s, a satellite link was established between SLAC and Argonne National Laboratory (ANL) to support an 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 encountering 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 Energy Research (ER) 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 with a proposal to upgrade the MFEnet.

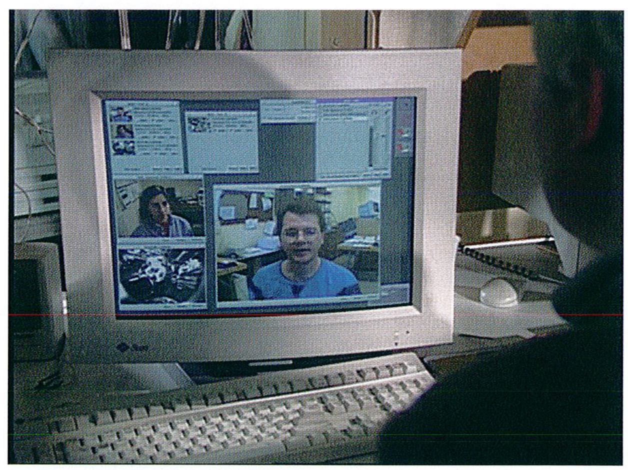

Later in FY 1985, Dr. Alvin Trivelpiece, then Director of Energy Research, charged ER's Scientific Computing Staff (now the Mathematical, Information, and Computational Sciences Division or MICS) with surveying computer networking requirements across all the ER programs and evaluating the status of existing 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 the ER programs, and to expand access to existing supercomputer facilities.

Dedicated Bandwidth for Collaboration Is Goal of R&D Project

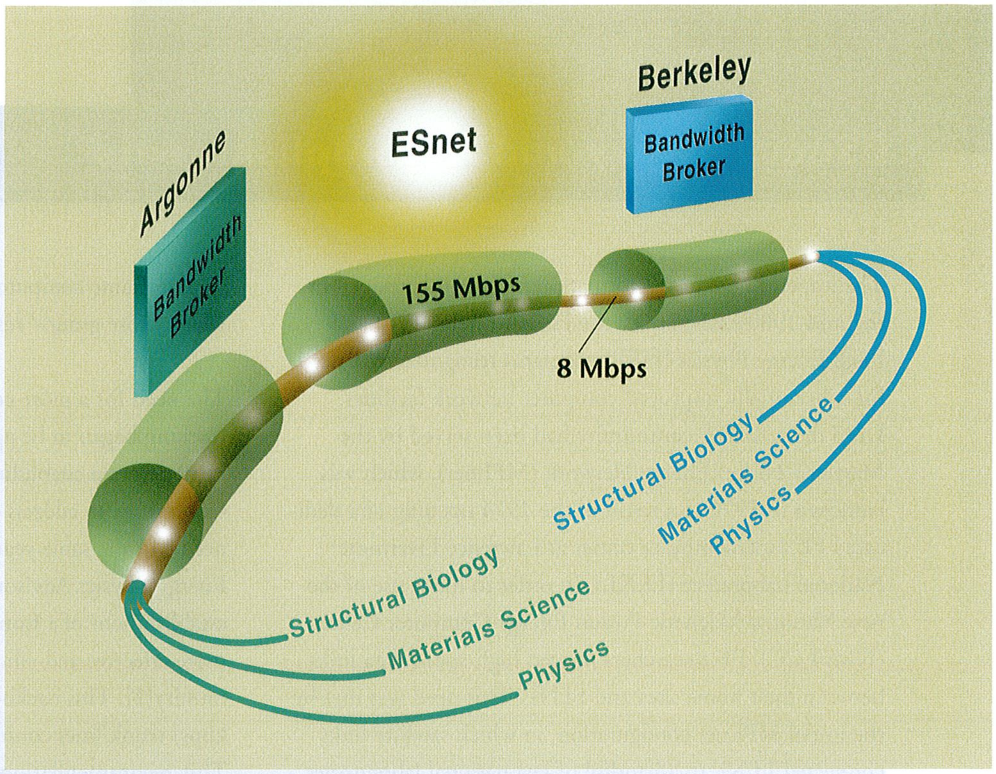

In the near future, bandwidth broker software at ESnet sites such as Argonne and Lawrence Berkeley national laboratories will guarantee as much as 8 megabits per second of bandwidth to three collaborative research projects. Berkeley Lab's Network Research Group is using DOE2000 funding to make this hypothetical scenario a reality.

Traffic jams on the Internet pose a significant problem for scientific collaborators. Real-time remote control of experimental facilities or computational simulations requires reliable quality of service (QoS)—in other words, guaranteed bandwidth, regardless of Internet traffic.

While major router vendors are starting to deploy the low-level machinery needed for QoS, the current standard for using this machinery to create end-to-end QoS is inadequate for collaborative research. There are two core problems:

- Traffic flow follows network topology, but bandwidth allocation decisions must follow organizational hierarchy. Researchers need separate, organizationally based agents to control bandwidth allocation.
- Allocation decisions involve both short-term considerations (particular conversations, data transfers, experiments, simulations, etc.) and long-term considerations (provisioning, budgeting, agency and organization priorities). Deciding how much to allocate should be separate from deciding which collaborators should use the allocation at a given time.

The ESnet QoS R&D project has two objectives:

1. Supply useful, deployable, flexible QoS to the Energy Research community over ESnet and in interaction with other networks.
2. Demonstrate to industry how administratively heterogeneous QoS might be implemented and deployed.

These objectives will be reached in two steps:

1. Enable and configure class-based queuing (CBQ) machinery in ESnet Cisco routers and selected site routers. CBQ is a traffic management technique that classifies incoming traffic according to a hierarchical class structure and applies queue-based throughput criteria on a per-class basis according to parameters contained in the user's traffic management policy.
2. Write and deploy a bandwidth broker application to control access to the CBQ machinery, enabling high-priority applications to receive sustained bandwidth allocations.

By giving scientists capabilities that are not yet available on commercial networks, ESnet creates new opportunities for scientific research while contributing to the evolution of the Internet.

As a result of these findings, Dr. Trivelpiece recommended that the MFEnet and HEPnet initiatives be combined and broadened 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 MICS 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 (ESSC) to represent the ER scientific community. The MICS 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 ER Scientific Computing 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 NMFECC, 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.

ESSC 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 ER-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. In February 1995, ESnet was the first major network to convert to a new technology—Asynchronous Transfer Mode (ATM)—and today over 30 sites are directly connected at speeds as high as OC3 (155 Mbps).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, ESnet provides seamless, multiprotocol connectivity among a variety of scientific facilities and computing resources in support of collaborative research, both nationwide and internationally. 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 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 network 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 ESnet is to satisfy these needs as fully as possible for DOE researchers. The networking needs of individual research programs are detailed in the following sections of this Program Plan.

References

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

ESnet Core Programs and Requirements

Stanford Linear Accelerator Center

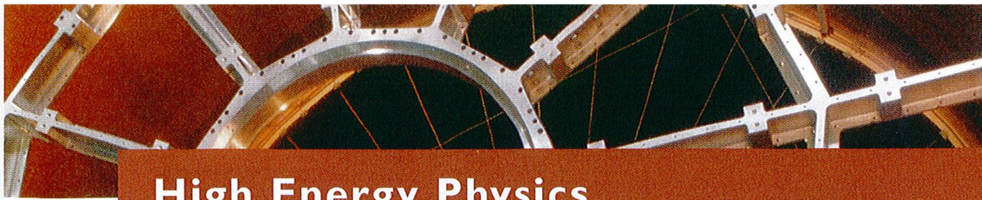

High Energy Physics

High energy physics (HEP) 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 (Fermilab), the Stanford Linear Accelerator Center (SLAC), Brookhaven National Laboratory (BNL), and Cornell University's Wilson Synchrotron. U.S. physicists are active users of accelerators abroad as well.

U.S. HEP researchers have initiated a major effort at the Large Hadron Collider (LHC), now under construction under an international agreement at CERN near Geneva, Switzerland. Among other goals, this facility will search for the Higgs boson. At the same time, the B Factory under construction at SLAC will be a significant new U.S. HEP site, and an international community of physicists interested in studying matter-antimatter asymmetry and charge-parity violation is completing work on a detector for that facility.

About 140 universities are an integral part of the U.S. HEP program. University physicists work on experiments at accelerator laboratories both in the U.S. and abroad, and on a variety of non-accelerator experiments. Networking is particularly vital to the HEP program because of the large international collaborations and the need for close collaboration between laboratory and university groups.

Major Components of the HEP Program

The U.S. HEP program is supported by both the U.S. Department of Energy (DOE) and the National Science Foundation (NSF). DOE operates several large facilities and supports the research of many university groups. The NSF operates one facility and supports university groups. The U.S. accelerator laboratories include:

- Fermi National Accelerator Laboratory, operating the 1.8 TeV Tevatron antiproton-proton collider and an 800 GeV fixed-target program. A major upgrade to the collider injector (Main Injector) will permit the collider program to maintain its vitality as the highest-energy collider in the world for several more years.
- The Stanford Linear Accelerator Center, operating the 91 GeV electron-positron linear collider (Stanford Linear Collider or SLC) and a fixed-target program at electron energies up to 50 GeV. An electron-positron "B-factory" (PEP-II) is under construction at SLAC, along with a major detector (BABAR).
- Brookhaven National Laboratory, operating the 30 GeV high-intensity Alternating Gradient Synchrotron (AGS) proton accelerator. Fermilab, SLAC, and BNL accelerators are all funded through the HEP program at DOE.
- The Cornell Electron Storage Ring (CESR), which provides electron-positron collisions at about 10 GeV, is funded by NSF.

These accelerators provide the core facilities with which the U.S. HEP research program is conducted, and attract scientists from around the world to participate. Two laboratories that operated accelerator facilities in the past, Lawrence Berkeley National Laboratory (LBNL) and Argonne National Laboratory (ANL), have a large technical infrastructure that provides critical support for the program.

U.S. physicists also participate in experiments at accelerator laboratories abroad:

- CERN (European Organization for Nuclear Research, Geneva, Switzerland), with both a 180 GeV Large Electron-Positron Collider (LEP) and a fixed-target program in operation, and a 14 TeV proton-proton collider (Large Hadron Collider or LHC) under construction
- DESY (German Electron Synchrotron Laboratory, Hamburg, Germany), with a 300 GeV electron (positron)-proton collider (Hadron Elektron Ring Anlage or HERA)
- KEK (Japanese National Laboratory for High Energy Physics, Tsukuba, Japan), with a broad program involving electron and proton beams at energies up to 25 GeV
- BEPC (Beijing Electron-Positron Collider, China) with an electron-positron collider at 2-5 GeV.

Large Hadron Collider Probes Mysteries of Matter

The Standard Model of fundamental particles and interactions has been remarkably successful, but it still leaves many questions unanswered: Why are there three types of quarks and leptons of each charge? Is there some pattern to their masses? Are there more types of particles and forces to be discovered? Are the quarks and leptons really fundamental, or do they too have substructures? Which particles make up the dark matter in the universe? How can gravitational interactions be included in the Standard Model?

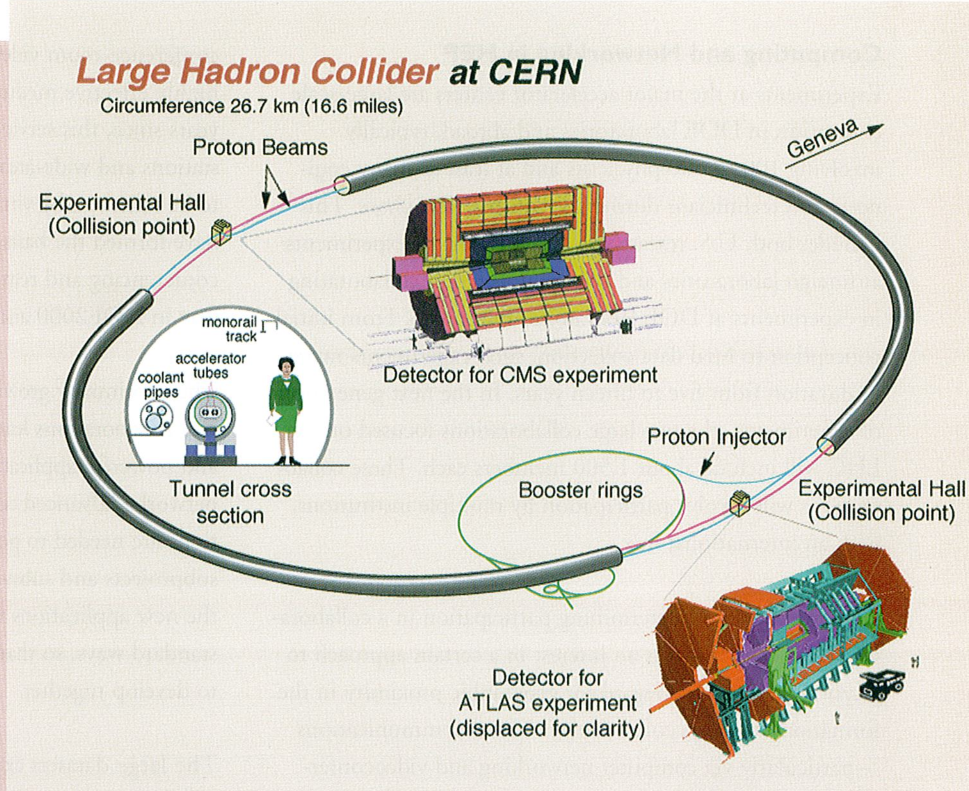

By accelerating each of two counter-rotating beams of protons to 7 trillion electron volts per proton (99.9999991% of the speed of light), the Large Hadron Collider will provide some missing pieces to the puzzle of the long-sought Grand Unified Theory.

Questions like these drive particle physicists to build new accelerators so that higher-energy collisions can provide clues to their answers. American physicists are collaborating on the construction and operation of a new collider at CERN on the Swiss-French border that will create collisions with the energy of 14 trillion electron volts—seven times higher than currently available. The new machine, called the Large Hadron Collider (LHC), will be turned on for data taking in 2005. With its high energy and the longest tunnel of any particle accelerator in the world—16 miles—the LHC will enable scientists to penetrate further into the structure of matter and recreate the conditions prevailing in the universe just seconds after the Big Bang.

One question that the LHC was specifically designed to answer is the mechanism by which electromagnetic and weak interactions become differentiated at low energies, although they are unified at high energies—a phenomenon called electroweak symmetry breaking. Speculations about the mechanism range from a single new particle, known as the Higgs particle, to a new dynamics, one example of which goes by the name supersymmetry. Physicists also hope that higher-energy data will provide clues about the multiplicity of ad hoc parameters in the current theory, including the masses of each of the elementary particles.

Teams in the U.S. are participating in the building of the collider itself and also in the construction and use of the two largest detectors. These two detectors, optimized for observation of hard collisions that produce previously unseen particles, have the names ATLAS and CMS. Each of these detectors has the participation of around 30 U.S. university and laboratory groups, along with over 100 other institutional groups from Europe and the rest of the world, totaling 1700 physicists from 33 countries. The networking services of ESnet are indispensable for a project of this size and complexity.

Computing and Networking in HEP

Experiments at the major accelerator centers are large-scale enterprises in DOE laboratories and abroad, typically involving 100 to 500 physicists and at least as many engineers and technicians during the construction phase. This includes both U.S. researchers collaborating in experiments at foreign laboratories and foreign researchers collaborating in experiments at DOE laboratories in the U.S. 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 1,500 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 computer 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 basic teleconferencing must operate with complete reliability and at high speeds.

In 1990, HEP groups began an experiment in the use of videoconferencing for scientific collaborations. With an initial link established between LBNL and SLAC, conference-room videoconferencing was found to be a highly effective medium for collaborative meetings. In the years since, this service, and the related service using workstations and wide-area network infrastructure, have grown to be relied on by virtually all programs using ESnet. They have formed the basis for the ambitious developments in conferencing and remote operation of facilities now under way in DOE2000 and other programs.

The continually growing complexity of HEP experiments and collaborations leads to a requirement for more advanced and powerful applications to enable collaboration via the network. Advanced conferencing and telepresence applications are needed to permit efficient joint work on even the subprojects and subsystems of large detectors. In general, the new applications will use the network in new and nonstandard ways, so that applications and networks will need to develop together.

The large datasets expected from Tevatron and PEP-II experiments (100 terabytes per year) in the near future and LHC experiments (1 petabyte per year) a few years later will create the need for close integration between the network infrastructure and the database applications that provide access to data for physics analysis purposes. Most experiments envision a system of regional centers that provide local access to the large datasets. Although the base datasets might not be transmitted over the network, and anyway are not subject to rapid change, smaller, partially processed "summaries" that can be used for detailed physics analysis by distributed analysis groups will need to be transmitted and cached in sophisticated ways that anticipate usage by physicists.

Although not an exhaustive list, the two new or improved applications listed above—telepresence and database access —are expected to be most needed by high energy physics.

Connectivity (Domestic and International)

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

B Factory Investigates Evolution of the Universe

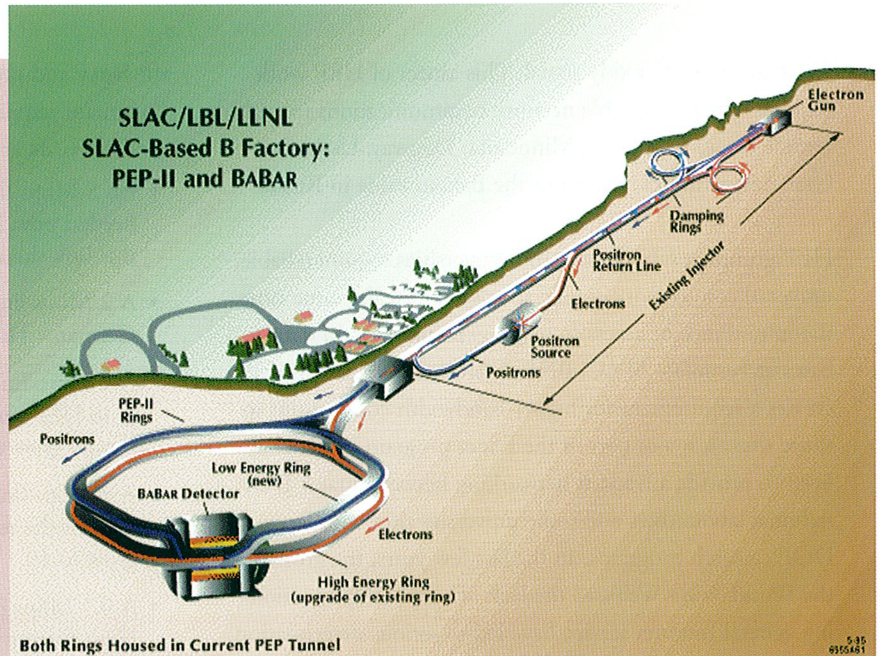

Billions of B mesons and anti-B mesons will be created at the SLAC B Factory. By examining their decay patterns, physicists hope to discover the mechanism by which the universe favors matter over anti-matter.

Why do we live in a matter-dominated universe? In the Big Bang, matter and anti-matter were produced in equal amounts, so what tipped the balance in favor of matter? This missing mechanism is crucial for a full understanding of the evolution of the universe. In primordial times there must have been a small imbalance in the matter-to-anti-matter ratio. All the matter and anti-matter, except the tiny unpaired excess, annihilated to form photons, which we can detect now. The unpaired excess—all matter—evolved to form the universe. To substantiate this model, we need experimental evidence that conclusively establishes the mechanism responsible for creating the tiny excess of matter.

To obtain this evidence, a new particle collider, Positron-Electron Project II (PEP-II), and a particle detector, BABAR (B/B-bar), will begin operating at the Stanford Linear Accelerator Center (SLAC) in 1999. The PEP-II accelerator will collide electrons and positrons to produce pairs of particles called B mesons and anti-B mesons, which contain b-quarks and anti-quarks. These particles will be detected by BABAR, which is being built by an international collaboration of 400 scientists from over 75 universities and laboratories in the U.S., Canada, Great Britain, France, Russia, Germany, Italy, Norway, and China. Because it will produce large quantities of B mesons, PEP-II has been nicknamed the "B Factory."

The PEP-II collider is an upgrade of the PEP facility that supported high-energy experiments in the 1980s, culminating in the first measurement of the b-quark lifetime. It will consist of two independent storage rings, one located atop the other in the PEP tunnel—the existing high-energy ring, and a newly constructed low-energy ring. Injection will be achieved by extracting electrons and positrons at collision energies from the SLAC linear collider (SLC) and transporting them in dedicated bypass lines.

Networking is extremely important for the B Factory's collaborative research. The BABAR detector will generate large amounts of raw data, from which the details of the beam collisions will be reconstructed. Monte Carlo programs will generate simulated data samples to calibrate the detector. The amount of raw, reconstructed, and Monte Carlo data is expected to be about 100 terabytes per year. Data will be sent over ESnet to institutions around the world, allowing prompt analysis and publication of results. Researchers already depend on extensive use of email, remote file sharing, collaborative code development and distribution, the World Wide Web, and videoconferencing. Some remote control of detector operations is also expected for problem diagnosis and debugging.

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 IHEP creates a requirement for higher-bandwidth connections to those sites. A major part of the ESnet program to date has been to provide advanced networking between major DOE national laboratories and some university sites with large DOE research programs. In the last few years, the growth of the World Wide Web and the lack of a smoothly functioning general Internet service have led to serious deterioration of connectivity between the DOE university researchers who are not directly connected to ESnet and the DOE national laboratories. New initiatives such as the very-high-speed Backbone Network Service (vBNS), the Next Generation Internet (NGI), and Internet 2 will attempt to improve the connectivity much needed by the university researchers, but ESnet will need to provide assistance when these efforts do not meet specific needs. 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. Major requirements of the HEP domestic program are:

- Continue to upgrade and strengthen connectivity between major high energy physics labs and other sites.
- Continue to monitor network performance between high energy physics researchers in DOE labs, foreign labs, and universities.
- Continue assisting in solving networking problems between high energy physics labs, universities, and foreign labs.
- Coordinate connectivity between ESnet and future domestic networks to optimize the networking required by the university high energy physics researchers to reach the DOE laboratories.

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 next five years are:

- An 8 Mbps link to CERN in Geneva, Switzerland, is needed now. As the LHC program becomes established, this link should grow steadily in capacity to 45 Mbps.
- A 2 Mbps link is needed to DESY in Hamburg, Germany. This link should grow in capacity to 10 Mbps.
- A 4 Mbps link is currently needed for general connectivity to HEP sites in the rest of Europe. This link's capacity should grow to 10.
- A 2 Mbps link is now needed to KEK in Tsukuba, Japan (and to the rest of Japan). This link's capacity should grow to 10.
- A 0.5 Mbps link is now needed now to the Institute of High Energy Physics (IHEP), in Beijing, China, with connectivity to other HEP institutes in China. This link's capacity should grow to 1.5 Mbps.
- 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 4 Mbps.
- A 1 Mbps connection is now needed to the Joint Institute for Nuclear Research (JINR) in Dubna, the Institute for High Energy Physics (IHEP) in 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.

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.

Nuclear Physics

Program Description

The scientific questions that the nuclear physics program addresses can be divided into four main areas[1]:

1. **Nuclear Structure and Dynamics:** This area investigates the properties and dynamic behavior of atomic nuclei, their production and decay patterns, and their behavior under extreme conditions (e.g., under rotations).
2. **Quark Structure of Matter:** We know that protons and neutrons, which are the constituents of the atomic nucleus, are in turn composite particles of quarks and gluons, which interact with each other in a very complicated way. Their behavior is studied using electron, proton, and photon beams, which are produced by particle accelerators.
3. **Phases of Nuclear Matter:** Whereas quarks and gluons are normally confined within the proton and neutron, theory no longer predicts this to be true under extreme conditions (so extreme as to exist only at the time of the Big Bang, when the universe was created). This newly predicted form of matter, called the quark-gluon plasma, is the subject of intense research using collisions of atomic nuclei at high energies (so-called relativistic heavy-ion collisions).
4. **Fundamental Symmetries and Nuclear Astrophysics:** Precision experiments at low and intermediate energies investigate the symmetries of nature at the atomic level. These symmetries have, among other effects, a significant impact on our understanding of the evolution of stars as well as the universe.

The list above demonstrates that nuclear physics is a very rich and diverse field. Advances have consequences and benefits beyond the narrow definition of the field. For example, improving our understanding of atomic nuclei and their decay patterns has an impact on our understanding of stellar evolution and applications as diverse as medical imaging, age determination in archeology, and ion implantation techniques to produce very hard surfaces for machining tools.

Service Requirements

Historically, nuclear physics has been performed by small groups of researchers located at universities or national laboratories operated by DOE or NSF. The experimental equipment was modest, and the amount of data taken and analyzed was small.

This picture has changed dramatically over the past ten years. While some experiments remain small, nuclear physics has seen a dramatic concentration of resources. The typical research project performs experiments at a few selected sites (mostly national laboratories but also some universities) that have the resources to operate medium to large accelerators, such as the Argonne Tandem Linac Accelerator (ATLAS) at Argonne National Laboratory, the Bates Linear Accelerator at MIT, the Holifield Radioactive Ion Beam Facility at Oak Ridge National Laboratory, and a few others.

The latest round of concentration has also seen the construction of the Thomas Jefferson National Accelerator Facility (JLab, formerly the Continuous Electron Beam Accelerator Facility, CEBAF). JLab has three experimental halls where 620 scientists from 120 participating institutions in 20 different countries perform their experiments.

An even bigger facility, the Relativistic Heavy Ion Collider (RHIC), is currently under construction at BNL, scheduled for completion in the spring of 1999. In four experiments at this facility, over 1,000 scientists will study heavy ion collisions at the highest available energy in the world.

This trend toward large experiments done by collaborations of hundreds of scientists at a few selected sites dominates the networking needs of the nuclear physics community. Reliable networking connections at the highest possible bandwidth are essential for the seamless operation of experiments of such magnitude, particularly with researchers located around the world.

Connectivity and bandwidth are clearly the two most important services provided by ESnet. However, other networking needs are becoming increasingly important, such as videoconferencing and directory services, which are very useful in large collaborations with multiple sites.

Connectivity Requirements

Good connectivity between sites where the experiments are performed and where most of the experimental data are analyzed is of premier importance to the nuclear physics program. Most of these sites are currently national laboratories with direct access to the ESnet backbone. The quality and reliability of these connections is vital for the success of the nuclear physics program. It should be noted that the NSF also supports two nuclear physics facilities: the Indiana University Cyclotron Facility and the National Super-

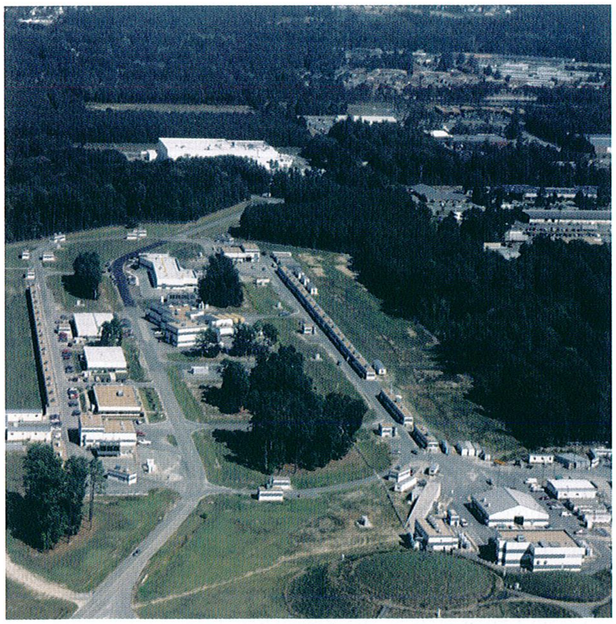

The Thomas Jefferson National Accelerator Facility (JLab, formerly CEBAF) in Newport News, Virginia, was built exclusively to study electron scattering in order to gain an understanding of the quark structure of matter. By 1998 the laboratory expects to produce 1 terabyte of data per day. Most of the data will be reconstructed and reduced to smaller data samples on site. The final analyses will be done at all the participating institutions, which therefore need good network access to the laboratory.

conducting Cyclotron Laboratory at Michigan State University. Hence, good connectivity between ESnet and the NSF-supported network is important.

Many researchers reside at universities and spend only a fraction of their time at the experimental sites. Good connectivity to ESnet is essential for them to successfully contribute to experiments that may be located over a thousand miles away. This can be achieved either through a direct connection of their computing facilities to ESnet or a guaranteed-efficient university-NSF-ESnet connection.

The interdisciplinary character of nuclear physics research also leads to additional connectivity requirements. For example, the joint Origins initiative, which supports research to understand the origins of the universe, spans three government agencies—DOE, NSF, and NASA.

International connectivity is also of great importance. Nuclear physics has seen international contributions increase substantially as the experiments and collaborations grow larger. In particular, JLab and RHIC have substantial contributions from European and Asian collaborators (mainly from Japan and China).

Capacity Requirements

Good networking bandwidth is needed in two respects: instant/interactive access and good average bandwidth.

Working interactively at an experimental site from a remote node over the network is essential for a large distributed collaboration. For example, the detector for the PHENIX experiment at RHIC consists of 11 different subsystems being built on three continents. Its readout electronics and trigger system have been designed and built in various locations across the United States. Most of the hardware designers will only spend limited time at Brookhaven, where the experiment will take place. It is thus important for them to be able to debug and monitor their hardware and software remotely.

However, not all required network activity is interactive. As outlined above, the experiments at JLab and RHIC will produce hundreds of terabytes of data per year, with an increase by a factor of 10 as RHIC reaches its design luminosity. Most of this data will be reduced to smaller data sets (order of several gigabytes) at a few selected sites where sufficient computing power is available. Researchers located at remote sites (mostly universities) subsequently need access to these smaller data sets so that they can analyze them at their home institution. A sizeable fraction of that data will be transferred over the network at a lower priority (e.g., overnight) than when instant interactive response is required.

Nuclear Physicists Seek Elusive Quark-Gluon Plasma

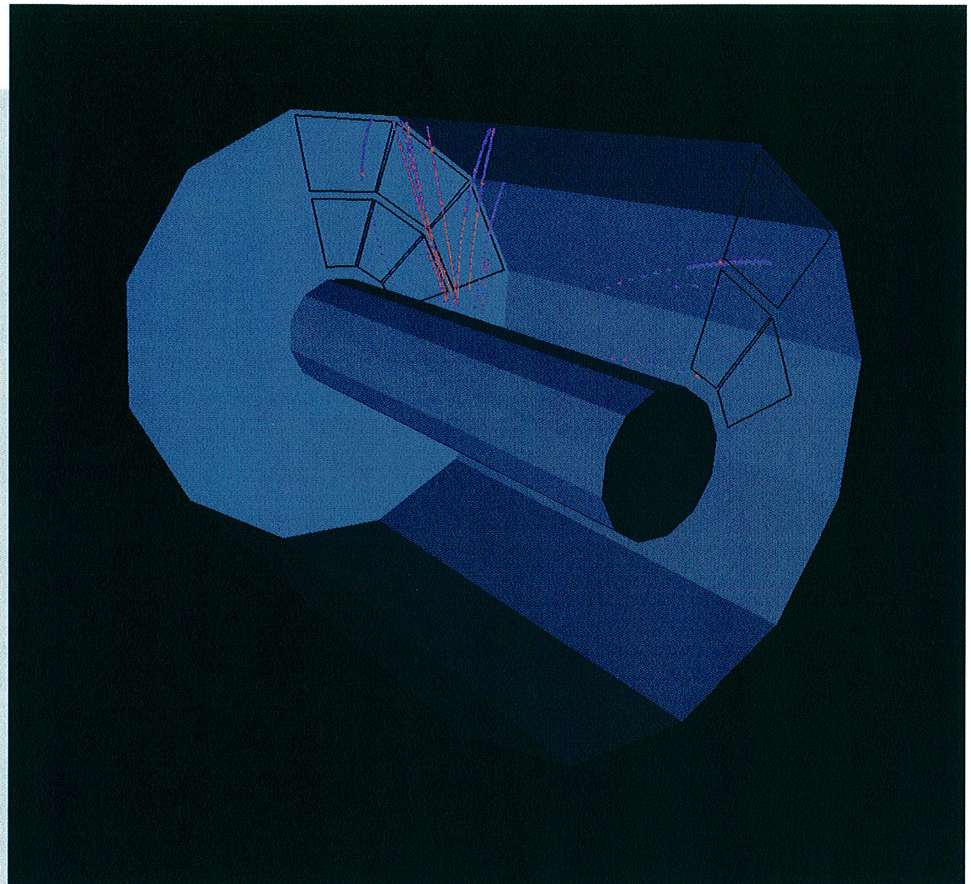

The STAR detector will record the locations of ionization events (blue or red squares in this simulation) as the particles given off by ion collisions pass through the gas filling the chamber. The events will then be reconstructed into tracks (red lines) by pattern recognition software.

Creating and detecting a state of matter called the quark-gluon plasma, which hasn't existed in the universe since a few microseconds after the Big Bang, is the quest of a new experimental physics facility at Brookhaven National Laboratory.

The Relativistic Heavy Ion Collider (RHIC) is scheduled for commissioning in 1999. Once completed, RHIC will counter-rotate beams of ions ranging from gold to hydrogen (protons) and collide them at a top energy of 100 GeV per nucleon for gold and 250 GeV for protons.

The goal of the RHIC project is to study the behavior of nuclear matter in the high-density, high-temperature conditions of these collisions. The theory of quantum chromodynamics predicts that the nuclear matter will undergo a phase transition from ordinary matter to the quark-gluon plasma. The experiments may confirm the current theory or reveal unanticipated wonders.

Four detectors (called BRAHMS, PHENIX, PHOBOS and STAR) are being built to measure the ion collisions and search for the quark-gluon plasma. They are designed to be sensitive to different signatures of the quark-gluon plasma, and together they will provide a detailed picture of the properties and behavior of nuclear matter under the extreme conditions of these collisions.

The STAR detector will measure and correlate many observables, such as particle tracking and momentum analysis, on an event-by-event basis. It will measure and identify virtually all charged particles over two units of rapidity (centered at mid-rapidity), a formidable computing task because hundreds of particles will be produced per collision. STAR will characterize particles at high transverse momenta and jets that are the result of hard scattered partons.

Analysis of data from RHIC and other next-generation physics experiments will require access to petabytes of data by hundreds of teraflops of processing power. Without the infrastructure provided by ESnet, management of this data would be a forbidding task.

Collaboratory Requirements

None of the existing major nuclear physics experiments are currently involved in any of the collaboratory efforts of the DOE2000 initiative. However, deliverables such as electronic logbooks and collaborative session management will certainly find application within the nuclear physics community as they become available to the researchers.

The Relativistic Heavy Ion Collider (RHIC), currently under construction at Brookhaven National Laboratory on Long Island, New York, is scheduled to be completed in spring 1999. Four experiments will observe heavy ion collisions in order to find the quark-gluon plasma and measure its properties. Data will be recorded at 50 megabytes per second. The expected data volume is 1.5 petabytes (1,500,000 gigabytes) per year.

Grand Challenge Requirements

The problems associated with handling the unprecedented amount of data produced by the next generation of high energy and nuclear physics experiments are the subject of one of the Grand Challenge applications (Data Access and Analysis of Massive Datasets for High Energy and Nuclear Physics). This Grand Challenge addresses the following three topics:

1. organization of data storage for efficient data management and selection of interesting events in a timely manner
2. providing the CPU power for the analysis and simulations
3. development of the software for selecting small data sets for subsequent analysis.

Another issue that this Grand Challenge will address is networking access for hundreds of users as well as transfer of smaller data sets (order of gigabytes) over ESnet. Data transfers of roughly 1 terabyte per week between RHIC and LBNL are expected by the end of 1998.

References

1. The DOE/NSF Nuclear Science Advisory Committee. Nuclear Science: A Long Range Plan. February 1996. http://pubweb.bnl.gov/~nsac/

Fusion Energy Sciences

Program Description

The mission of the Fusion Energy Sciences (FES) Program is to provide a science base for the development of 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. Experimental research in this field requires the acquisition, analysis, and visualization of very large numbers of multidimensional signals that describe plasma parameters as functions of space, energy, and time.

In this country, FES work is distributed across many 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 FES 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. The current generation of U.S. experiments are explicitly national facilities, with advisory committees made up of experts from outside laboratories and with a host of collaborators participating in their experimental programs. The next large experiment, the International Thermonuclear Experimental Reactor (ITER)[1, 2], is being designed by an international team spread over three continents.

Theoretical research in FES has become more and more collaborative as the complexity of the issues to be addressed has increased. Typically, national and international teams or expert groups have been formed to tackle issues of importance to ITER. Major code development has also become a collaborative enterprise under the umbrella of such national projects as NIMROD and the Numerical Tokamak Turbulence Project (NTTP)[3] because of the complexity of the models and the sophistication of their computational implementation. In addition, the ties to theoretical efforts in Europe and Japan have been strengthened through working groups of the Transport Task Force or by the establishment of projects of a similar scope and focus abroad, such as the numerical tokamak initiatives in Europe and Japan. Remote collaboration is also the preferred mode of operation for validation of theoretical calculations with experiments performed at major U.S. experimental sites.

Service Requirements

Fusion energy research requires all basic network services, including electronic mail, file transfer, remote log-in, remote printing and plotting, directory services, and World Wide Web access. In recent years, however, remote experimental and theoretical collaborations have added support requirements beyond these "lifeline" services. Computing, visualization, and file management for a given task are already shared over the network. A large number of FES 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, FES researchers will soon need such capabilities as remote or distributed code management and distributed system management. Use of the network for access to experimental equipment has heightened the importance of security—already important for any enterprise dependent on its computers and networks. Shared security infrastructure must be a part of our future environment.

Because current design activities and experimental and theoretical collaborations are increasingly multi-institutional as well as multinational in scope, the FES community will require more powerful tools for interpersonal communications, such as videoconferencing and multimedia mail. It is now assumed that the ITER program will set up distributed control rooms to support data acquisition and instrument

Fusion Research Tries to Tame Immense Forces

Nuclear fusion, the force that powers the sun and stars, has the potential to be an abundant source of energy for the future. But designing a system that will contain and control such a turbulent force is a major challenge. To harness fusion as a practical energy source, the fuel must be heated to more than 100 million degrees and confined long enough to be burned efficiently.

With magnetic confinement, the hot ionized fuel is contained by powerful magnetic fields, and the fuel burns continuously or in pulses lasting hundreds or thousands of seconds. In a different approach, inertial confinement, the fuel is heated and compressed by intense laser or ion beams; the plasma is effectively confined by its own inertia, and burning occurs in less than a nanosecond.

The leading design for magnetic confinement is a toroidal device called a tokamak. From their invention in the Soviet Union in the 1950s, tokamaks have progressed to "scientific break even," the point where the heating power applied is balanced by the fusion power produced. Fusion yields per pulse from these experiments have increased by a factor of a trillion over the last 25 years, exceeding the rate of improvement in the semiconductor industry.

Fusion research is highly collaborative, with researchers from dozens of labs and universities carrying out experiments. Beginning in the late 1980s, scientists have made measurements from remote sites, using ESnet to control plasma diagnostics, view data, and interact with their colleagues via audio and video. More recently, major tokamak facilities have been controlled entirely from offsite, with a remote team running the experiment and directing a team of researchers spread around the country.

Research conducted at experimental facilities is complemented by the Numerical Tokamak Turbulence Project (NTTP). The primary research objective of NTTP is to develop a predictive ability in modeling turbulent transport due to drift-type instabilities in the core of tokamak fusion experiments. This is accomplished through the use of three-dimensional kinetic and fluid simulations and the derivation of reduced models.

The NTTP simulations build a bridge between theory and experiment by providing a deeper understanding of experimental results and design options. Since controlling the energy transport has significant leverage on the performance, size, and cost of fusion experiments, reliable NTTP simulations can lead to significant cost savings and improved performance in future experiments. ESnet plays an essential role by providing access to the high-performance computers required for these simulations.

NTTP researchers completed full cross-section fluid calculations of cylindrical multihelicity ion-temperature-gradient-driven turbulence with previously unattainable resolution on the Cray T3E-600 at the National Energy Research Scientific Computing Center (NERSC). Work is in progress to increase the resolution, improve the performance of the parallel code, and include toroidal geometry in calculations on NERSC's T3E-900.

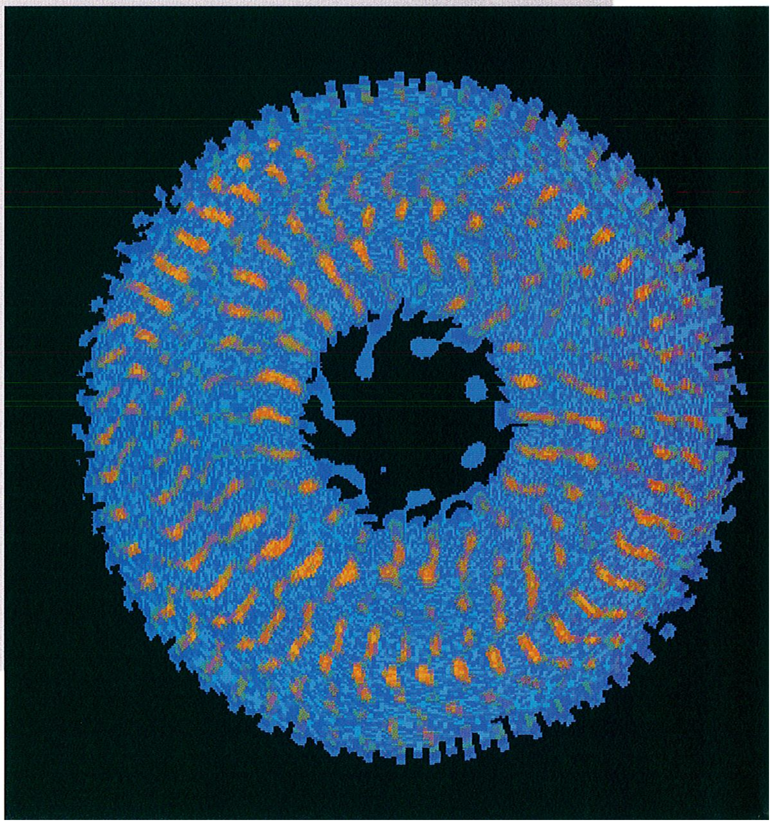

control for experimenters located at remote sites. The nature of the experimental methods employed in FES research requires close communication between all members of the physics teams responsible for acquiring and analyzing experimental data. While it is to be expected that many members of each lab's 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.

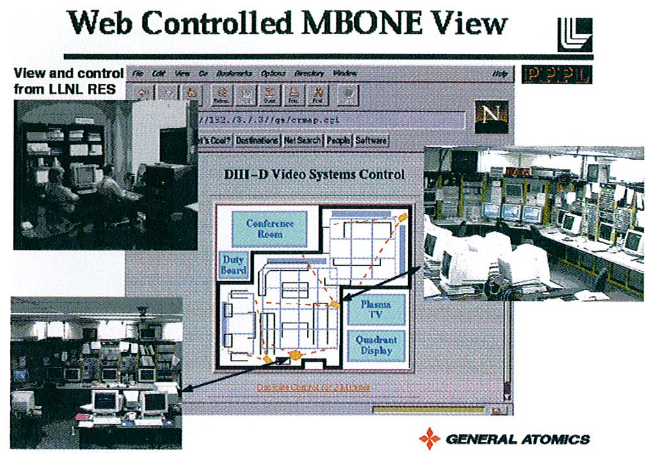

The Remote Experimental Environment (REE) uses Mbone over ESnet to communicate between multiple remote sites during magnetic fusion experiments at General Atomics, San Diego, California.

Connectivity Requirements

Worldwide sites involved in FES research include some 100 in the U.S., 65 in Europe, 45 in Japan, and several in South America, China, and Australia. For U.S. FES research sites, the highest priority for network connectivity is access to NERSC and to the major experimental sites at General Atomics (GA), MIT, and Princeton Plasma Physics Laboratory (PPPL). As theoretical and experimental collaborations have grown in size and importance, first-rate connectivity between all participating labs and universities has become essential.

FES 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 FES as our domestic connections are today.

Capacity Requirements

Today's ESnet, with its T3/OC3 backbone and its gateways to the broader Internet, generally meets the current domestic needs of the FES program. The most notable unmet needs are for high-speed tail circuits to non-backbone sites, particularly universities, and the lack of adequate high-speed links to most international sites. The processes described above (visualization, distributed computing, remote control of experiments, etc.) entail moving complex computing applications from local-area (LAN) to wide-area networks (WAN), and these procedures will require ever-increasing bandwidth. In this connection, it should be noted that most sites require FDDI and/or switched 100 Mbps Ethernet for their local infrastructure. We anticipate that requirements for the wide-area network will soon reach these levels.

Collaboratory Requirements

Researchers in the fusion energy program have begun creating a National Magnetic Fusion Collaboratory (NMFC)[4-6]. Our overall objective is to allow scientists at remote sites to participate as fully in our experiments and computational activities as if they were working at a common site. These objectives can be realized through a collaboratory that consists of four components:

1. remote experimental operations
2. remote code development
3. online code and data access
4. planning and coordination.

The primary benefit of the collaboratory to the fusion program is better support for ongoing and planned research. The collaboratory will enhance the effectiveness of remote participants, ease the requirements for frequent travel, and increase the number of researchers who can participate in experimental operations and analysis. The same environment will facilitate multi-institutional code development projects and improve the interaction between theory and experiment by providing greatly improved access to experimental data and modeling codes.

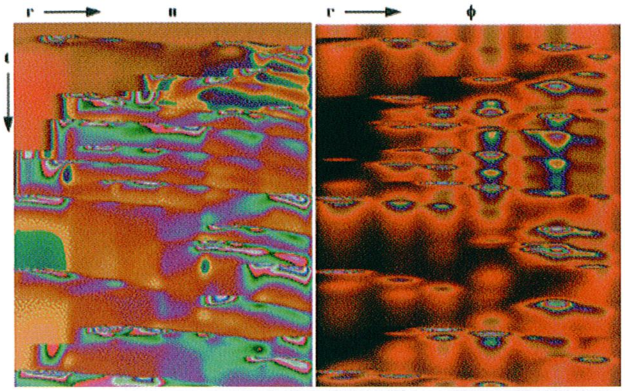

This simulation of self-organized criticality (SOC) dynamics for magnetically confined plasmas shows large transport events in the density on the left and smaller-scale fluctuations in the potential on the right, as expected from SOC systems. These large-scale 3D turbulence calculations were performed remotely on high-performance computers at NERSC.

The magnetic fusion research community has considerable experience in placing remote collaboration tools in the hands of real users. The ability to remotely view operations and to control selected instrumentation and analysis tasks has been demonstrated. University of Wisconsin scientists making turbulence measurements on the Tokamak Fusion Test Reactor (TFTR)[7] were provided with a remote control room from which they could operate their diagnostic while keeping in close contact with their colleagues in Princeton. LLNL has assembled a remote control room in Livermore in support of a large, long-term collaboration on the DIII-D tokamak at GA in San Diego[8,9]. From the same control room, a joint team of MIT and LLNL scientists has conducted full functional operation of the Alcator C-Mod tokamak located 3,000 miles away in Cambridge, Massachusetts[10]. These early efforts have been highly successful, but are only the first steps needed to demonstrate the technical feasibility of a complete "facilities on line" environment.

To aid distributed code development, teams will make use of desktop videoconferencing, a common code-version control system, an electronic notebook to document coding changes, common display output demonstrations, and other available communication and collaboration technologies. The NMFC will implement a common online code-sharing library to improve code and data access, code interconnection, and code invocation. This online facility will include codes for both modeling and data analysis. Construction of this resource is planned as a phased multi-year effort. In the first phase, I/O standards will be adopted and implemented to facilitate both the transfer of experimental data into the analysis codes and the transfer of processed output data among the codes. Network-based tools for code and database invocation and visualization will be developed. In subsequent phases, issues such as the efficient use of distributed compute cycles, reliable asynchronous intertask communications, multicasting data, remote display and downloading of results, distributed task queuing, and session management will be addressed. The development of this code-sharing facility will in itself be a demonstration project in remote collaboration, as it will be carried out by multi-institutional teams employing network collaboration tools.

Grand Challenge Requirements

The Numerical Tokamak Turbulence Project (NTTP)[3] is one of DOE's Phase II High Performance Computing (HPC) Grand Challenges (GC). It is aimed at an understanding of anomalous transport in magnetic confinement devices based on first-principles physics models implemented on massively parallel computers. The NTTP regroups fusion researchers, computational physicists, and computer scientists, located at ten geographically distributed institutions and three of the HPC Research Centers (HPCRCs), into a tightly knit, though remote, theoretical collaboration.

The GC requirements are complementary to those of the NMFC. These are driven by joint code development, access to geographically distributed computer resources, interactive steering of large-scale calculations, exchange of large amounts of data between participating researchers, and remote visualization of renditions of such data in four dimensions, all preferably in real time. The basic GC requirement is high bandwith and reliable connectivity between the participating sites, as well as the sites and the HPCRCs located at NERSC, the Advanced Computing Laboratory (ACL) at LANL, and the Center for Computational Sciences (CCS) at ORNL. The GC requirements encompass more than adequate bandwidth to accomodate high-resolution desktop audio and video tools in order to facilitate communications between the participants. They also include bandwidth on demand to enable, in production mode, virtual reality applications of the type demonstrated by the NTTP team as part of the I-WAY project at Supercomputing '95[11].

References

1 K. Tomabechi, J.R. Gilleland, Y. A. Sokolov, R. Toshci, and the ITER team, "ITER Concept Definition," *Nuclear Fusion*, 31 (6), 1135 (1991).

2 ITER Team, "ITER Interim Design Report," International Atomic Energy Agency, Vienna, 1996.

3 B.I. Cohen, J. M. Dawson, V. K. Decyk, W. D. Dorland, G. W. Hammett, G. D. Kerbel, J.-N. Leboeuf, W. W. Lee, J. V. W. Reynders, and R. E. Waltz, "The Numerical Tokamak Turbulence Project," an HPCC GCA proposal, June 19, 1996 (http://www.acl.lanl.gov/GrandChal/Tok/gcaproposal96).

4 D. Barnes, T. Casper, S. Davis, M. Greenwald, D. Greenwood, P. Henline, J. Moller, P. Roney, "Requirements for a National Fusion Collaboratory," white paper, October 1996 (http://www.pppl.gov/nfc).

5 B. Braams, J. Cary, V. Chan, R. Cohen, W. Dorland, W. Houlberg, S. Jardin, D. Schnack, "Needs and Opportunities in Computational Fusion Science," Report to OFES and MICS, October 1996 (http://cips.colorado.edu/NeedsOppsCompFusion.html).

6 J. Cary, T. Casper, V. Chan, J. Crotinger, S. Davis, M. Greenwald, D. Greenwood, S. Jardin, W. Houlberg, B. McHarg, "National Magnetic Fusion Collaboratory Proposal for a DOE2000 Pilot Project," 1996 (http://www.pppl.gov/nmfc/proposal.html).

7 R. J. Fonck et al., "Remote operation of the TFTR BES experiment from an off-site location," *Rev. Sci. Instrum.*, 63 (10), October 1992.

8 T. A. Casper and W. J. Lennon, "Remote Experimental Site: A Command and Analysis Center for Big Physics Experimentation," *Proceedings of the 14th IEEE/NPSS Symposium on Fusion Engineering*, San Diego, CA, October 1991.

9 W. H. Meyer, D. N. Butner, and T. A. Casper, Remote Experimental Site (RES) Release 1.0, *Bull. Am. Phys. Soc.*, 38, 2112 (1993).

10 S. Horne, M. Greenwald, T. Fredian, I. Hutchinson, B. Labombard, J. Stillerman, Y. Takase, S. Wolfe, T. Casper, D. Butner, W. Meyer, and J. Moller, "Remote Control of Alcator C-Mod from LLNL," *Fusion Technology*, 32 (1997), 52 (http://cmod2.pfc.mit.edu/~g/remote.html).

11 G. D. Kerbel, T. Pierce, J. L. Milovich, D. E. Shumaker, A. Verlo, R. E. Waltz, G. W. Hammett, M. A. Beer, W. Dorland, "Interactive Scientific Exploration of Gyrofluid Tokamak Turbulence," *International Journal of Supercomputing Applications and High Performance Computing*, 10, No. 2-3, 182-198, Summer/Fall 1996.

Basic Energy Sciences

Program Description

The Office of Basic Energy Sciences (BES) is responsible for basic research in the natural sciences, leading to new and improved energy technologies. BES supports fundamental research in areas related to energy resources, production, conversion, and efficiency, as well as areas related to mitigating the adverse impacts of energy production and use. As such, BES is the United States' foremost sponsor of fundamental research in broad areas of bioscience, chemical science, geoscience, materials science, and engineering science. The purpose of the program is to underpin DOE missions in energy and environment, to advance basic science on a broad front, and to provide unique national facilities for the scientific community.

To carry out its mission, BES funds research as programs within the DOE national laboratory complex and as research grants by principal investigators from universities and industry. BES also constructs and operates unique research facilities, such as synchrotron light sources for infrared, ultraviolet, and x-ray studies; neutron sources; and centers for electron beam microcharacterization, combustion research, materials preparation, surface studies, and ion beam studies. Each of these three kinds of BES programs has network service requirements of its own, as well as requirements common to one another.

BES User Facilities

- Advanced Light Source, Lawrence Berkeley National Laboratory
- Advanced Photon Source, Argonne National Laboratory
- Intense Pulsed Neutron Source, Argonne National Laboratory
- National Synchrotron Light Source, Brookhaven National Laboratory
- Los Alamos Neutron Scattering Center
- High Flux Isotope Reactor, Oak Ridge National Laboratory
- Stanford Synchrotron Radiation Laboratory, Stanford Linear Accelerator Center
- High Flux Beam Reactor, Brookhaven National Laboratory
- Materials Preparation Center, Ames Laboratory
- Electron Microscopy Center, Argonne National Laboratory
- Center for Microanalysis, University of Illinois
- National Center for Electron Microscopy, Lawrence Berkeley National Laboratory
- Shared Research Equipment Program, Oak Ridge National Laboratory
- Surface Modification and Characterization Research Center, Oak Ridge National Laboratory
- Combustion Research Facility, Sandia National Laboratory, Livermore
- James R. MacDonald Laboratory, Kansas State University
- Pulse Radiolysis Facility, Notre Dame University
- National Spallation Neutron Source (in advanced design phase), Oak Ridge National Laboratory

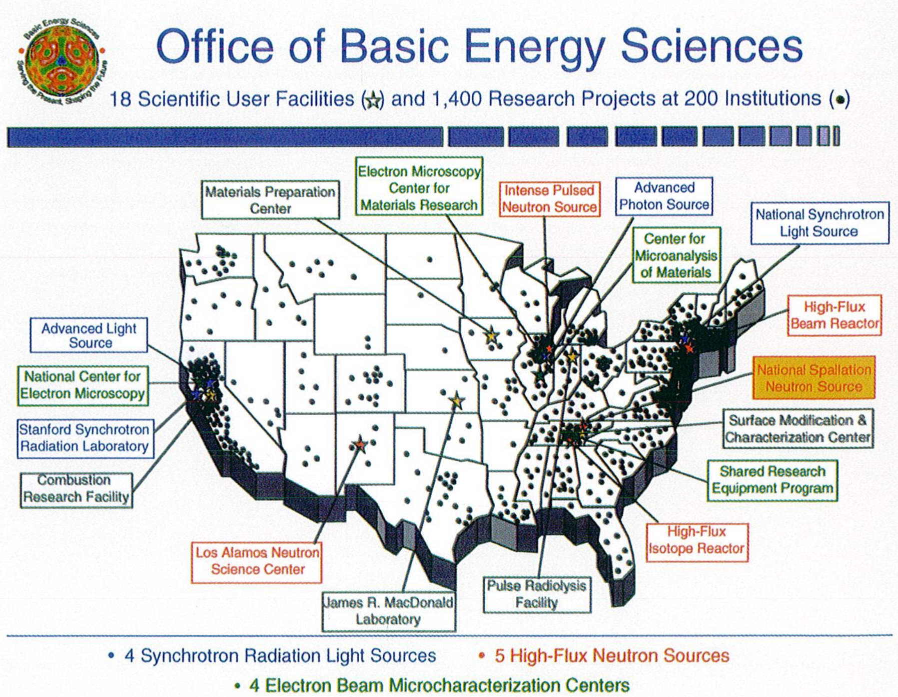

Service Requirements

The BES programs at DOE national laboratories are the focus of large, integrated, and often multidisciplinary efforts. In addition to on-site capabilities, BES researchers use experimental and computational facilities across the DOE research complex. A laboratory program will often involve a significant group within a local area network, coupled (increasingly) with collaborating groups at other laboratories, national facilities, or universities.

Research grants are representative of the most advanced frontier of fundamental understanding, and quite frequently are oriented toward professors, post-docs, or graduate students. Generally, grant participants are part of widely dispersed research communities.

The research facilities operated by BES are centers of expertise and instruments that provide unique capabilities to the national science community, including DOE, universities, and industry. Although many scientists travel to these facilities to use them, there is a growing requirement to operate instruments remotely whenever possible, and to participate remotely with onsite users. This offers the prospect of greater productivity and efficiency in the use of the facility, and also helps promote scientific and diversity goals (as facilities are more accessible). This can only succeed if the capabilities and information available to remote users are a significant fraction of what is available onsite.

The typical mode of operation at user facilities is to allot investigators a slot of time on the machine during which they collect as much data as possible. The data acquired are analyzed later. Some data analysis is done online, or at least locally, to guide the progress of the experiment and to reduce/compress the often voluminous raw data. The most efficient balance of local computing and distributed computing depends on the instrument and its support computers, and on the particular experiment. After the run, the data will either be analyzed at the facility or moved elsewhere for analysis. Data connections are one very important factor influencing these choices. During the analysis phase, collaborations must be maintained between the investigators themselves (many of whom are not DOE investigators but are participating as a consequence of the general service mission assigned the facilities) and also with the instrument scientist(s) located at the facility. Increasingly, research projects involve data obtained at more than one facility. Network connections are essential for connecting and coordinating these collaborations.

Large-scale computations involving modeling, simulation, and data analysis play a growing role in BES investigations. As the size and complexity of these computations and their data grow, network requirements also increase. In fact, limitations of current networks continue to constrain the use of high-performance computing facilities by putting significant bounds on how much information can be transported and where large data can be analyzed.

Collectively, the success of BES programs today and in the future is ever more closely tied to network services. High-performance networking is increasingly a critical success factor for computations, data sharing, remote instrument access, and collaboration. Network shared-file systems, directory services, authentication, and data security are critical infrastructure needed by BES programs. Finally, requirements for Internet videoconferencing and scientific collaboratories are emerging from a wide range of BES programs because of their power to coordinate and facilitate complex research programs.

Connectivity Requirements

The ideal model of network interconnections is that any data, no matter how big, could be transported from any point to any other point in the system quickly enough to be transparent to the researcher's work. The system, in the instance of Basic Energy Sciences, consists of all the facilities, the national laboratories, the institutions of the principal investigators, users of the facilities and their collaborators, and even the homes and field/mobile locations of all personnel involved. Though good science will arise from seeking this ideal, some practical priorities must be considered. In the real world, network bandwidth is a limited, expensive resource. Wise management requires that those resources be put where maximum-value applications will be enabled.

Connectivity to (and among) the DOE national laboratories and other BES user facility locations is paramount for BES. This entails connectivity to the widespread set of universities that use DOE facilities. It must also be recognized that not all components are within the purview of ESnet—and thus a key consideration is that ESnet enables, and perhaps even assists, other parts of the network in achieving the whole.

Reliability continues as an issue, and quality-of-service control is becoming essential. Many research activities today are designed to heavily exploit the connectivity of the network and cannot proceed without reliable connections. This will require both new services and ongoing improvements in coordinating trouble-tracking among ESnet and its peers.

Capacity Requirements

In addition to excellent management, ESnet has benefited from having a restricted customer base. Consequently, it has not suffered some of the degradation caused by over-subscription of other networks. However, larger-scale tasks, which have really only just begun to appear, will require 2 to 3 orders-of-magnitude increase in bandwidth. Large-scale computational efforts have historically led experimental efforts in requiring the ability to move large data sets quickly, and those efforts continue to grow (see Grand Challenge Requirements below). However, experimental research collaborations (e.g., at BES user facilities) have begun to compete very strongly with those demands, since the newer experimental facilities produce enormous quantities of data at prodigious rates. These data must be utilized in most cases through extensive, computationally demanding analysis and graphical representation. Shared analysis and transport of these data, often in real time, will grow rapidly. Also, large data sets must be moved from the experimental facility to a computational facility, and large data repositories must be made accessible to a broader community. Inevitably, network capacity requirements are growing very rapidly.

Collaboratory Requirements

Electronic collaboration has shown the promise of becoming a primary enabling factor in research, with a growing number of successful pilot projects in BES facilities. To achieve this "laboratory without walls" requires network infrastructure, network services, and tools from a number of sources. The broad scope of BES programs, the geographic distribution of BES resources, and the heterogeneous nature of computing equipment creates significant challenges for creating and deploying the needed capabilities. Interoperability (the ability of applications on different kinds of computers to work together) and extensibility (the ability to extend applications with new functions) are crucial. To achieve this, the DOE community needs to work together, as a community, leveraging and helping to define standards that will permit DOE and others to develop tools that can be integrated into the collaboration systems we need for our research. Next to bandwidth requirements, collaboratory capabilities are the most sought-after network resource requested by representatives of the BES community.

Remote monitoring or operation of BES experimental facilities has become feasible and is showing some excellent initial success. For example, the Spectro-Microscopy Collaboratory has provided investigators in Wisconsin and elsewhere with remote access to x-ray microscopy facilities at the Advanced Light Source in Berkeley. BES is also co-sponsoring two major pilot collaboratory projects in the DOE2000 program:

- The Materials Microcharacterization Collaboratory will integrate and coordinate efforts of electron microscopists across the country into a virtual laboratory for microcharacterization of interfaces and surfaces (http://tpm.amc.anl.gov/MMC/). Partners include Argonne, Lawrence Berkeley, and Oak Ridge national laboratories, the University of Illinois Frederick Seitz Laboratories, the National Institute of Standards and Technology, Gatan Inc., R. J. Lee Group, EMiSPEC Systems Inc., Philips Electronic Instruments, Hitachi Instruments, Inc., JEOL-USA Inc., Sun Microsystems, and Graham Technology Solutions.
- The Diesel Combustion Collaboratory is investigating the use of computer and networking technologies to improve communication between participants in the Heavy Duty Diesel Combustion CRADA (Cooperative Research and Development Agreement) (http://www-collab.ca.sandia.gov/Diesel/). Partners include Sandia, Lawrence Berkeley, Lawrence Livermore, and Los Alamos national laboratories, the University of Wisconsin-Madison, Cummins Engine Co., Caterpillar Inc., and Detroit Diesel.

For collaboratories to be successful, the presence of the investigator must be projected to the remote site, and the full feel of the site must be projected to the investigator. Not only must the investigator be able to observe what is going on, that investigator's presence must be felt at the facility site. The emerging paradigm makes heavy use of videoconferencing, shared screens and visualization, remote operation, and electronic notebooks, which require much wider bandwidth than a simple logon. Such connections have the additional complication of needing to go to wherever the investigator is located. Often that implies going through another network interconnected to ESnet.

Collaboratory Advances Materials Research

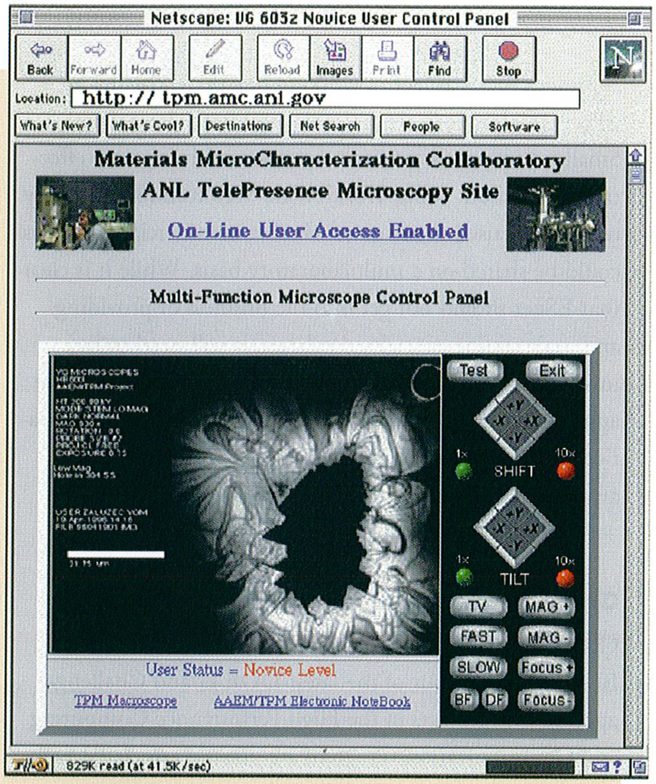

New collaboratory technologies allow Argonne National Laboratory's electron microscope to be controlled remotely over the World Wide Web.

Large-scale changes in production and distribution processes led to the first industrial revolution. Today, scientists are looking to create a revolution on a much smaller, even microscopic scale. In fact, by allowing scientists to study and create materials at near-atomic levels, ESnet is helping provide a closer understanding of why materials behave the way they do. These behaviors can produce such desirable properties as resistance to wear and breakage.

Technologically advanced materials exhibit unique properties because their microstructure and microchemistry are carefully tailored during the manufacturing process. However, the tools needed to better understand these microscopic features are expensive and available only at specialized research facilities.

The DOE2000 Materials MicroCharacterization Collaboratory (MMC) links scientists at five materials microcharacterization facilities in an interactive electronic laboratory. Participants include Argonne, Lawrence Berkeley, and Oak Ridge national laboratories, the National Institute of Standards and Technology (NIST), and the University of Illinois. Each of these facilities possesses unique, dedicated instrumentation not available elsewhere. ESnet provides the bridge that brings these institutions together and enables each to build upon the others' efforts.

MMC participating centers collectively house virtually every characterization technique which employs electrons, ions, photons, x-rays, neutrons, mechanical, and/or electromagnetic radiation to elucidate the microstructure of any material. MMC provides other participants with easy remote access to these complementary facilities via ESnet connections.

Telepresence operation of scientific instruments is a developing technology among MMC partners, and each center focuses on a particular aspect. Argonne is concentrating on platform-independent access and control, implemented via an intelligent electronic laboratory server that uses conventional World Wide Web browsers. Berkeley Lab is exploring an intelligent human-machine interface for remote dynamic experiments in which the Internet cannot guarantee real-time delivery of commands and data. Oak Ridge is focusing on automating as much of the microscope operation as possible on the local computer and then using low-cost commercial software to remotely view and control the local computer.

While access to instrumentation is important, new scientific discoveries depend even more on the collaboration of scientists who design and execute the experiments. By placing creative scientists with complementary expertise together in a new environment that allows convenient and dynamic interactions, unencumbered by distance, the MMC enhances their ability to conduct scientific research. Using ESnet's capabilities and services, the Collaboratory enables scientists to work virtually side by side, sharing data, observations, and expertise in real time. That's a real advance in how scientific research is carried out.

Another crucial aspect of working together, whether as direct collaborators or not, is the ability to share data. Small- to intermediate-sized databases are being very effectively served out using World Wide Web technology. Many users have raised the issue of whether commercial databases could be shared on a multilaboratory basis. While it is clear that ESnet should not get into the business of providing and maintaining specialty databases, it will need to be involved in providing the capability to deliver them to the appropriate users. Similarly, as users expand their collaborations, there will be a need to have large data collections available remotely.

Grand Challenge Requirements

Grand Challenge computing is an essential element of BES programs. Four of the latest DOE Grand Challenge applications projects in the High Performance Computing and Communications Program are directly associated with BES research efforts, and large-scale computing is an integral part of many BES efforts in chemical and materials science. These efforts not only use a widely distributed array of DOE's computing resources, they also feature widespread collaborations. Hence, today's Grand Challenge projects combine the challenging network requirements of high-performance computing with the network challenges of collaboratories.

For instance, Grand Challenge collaborations on software development are curtailed by problems in establishing the shared file systems needed to carry out the distributed effort. The difficulties stem from both insufficient network performance and the lack of common client software that can operate securely over wide-area networks.

There is a growing level of computational effort that performs different parts of the calculation on the computer best suited for that type of work. Consequently, the network must move the intermediate data from the computer that has completed a task on to the computer that will take the next step. Other scenarios employ the network to integrate a high-performance computer with a database at another location, or a running simulation with remote visualization equipment. Conceptually, this is a realization of the concept that the network serves as the actual backplane of a "meta-computer" made up of all the resources available in the network. The concept has a reasonably broad acceptance: it was proposed as a basis for planning by the Energy Research Supercomputer Users Group at the beginning of this decade, and it had a spectacularly successful demonstration in the Information Wide Area Year, or I-WAY, project associated with the Supercomputing '95 conference. Such approaches are computer efficient, but require much higher performance networks (into the gigabit-per-second range with very low latency), along with mechanisms to dynamically control network connection quality of service.

Mathematical, Information and Computational Sciences

The mission of the Mathematical, Information and Computational Sciences (MICS) Division is to contribute to DOE programs and to the national interest through research and the application of advanced mathematical, computational, computer, and communication sciences. These are accomplished through both the High Performance Computing and Communications Program and the National Information Infrastructure Initiative. MICS-managed efforts, categorized into five elements, are summarized here:

High Performance Computing Systems. This element is meant to ensure that future generations of hardware and software, including architecture, are useful throughout a broad range of DOE applications. Current investigations focus on the exploration of the current and future utility of hardware and software parallelism.

Advanced Software Technology and Algorithms. Through this element, high-performance computer systems are made available to significant Energy Research (ER) applications and national Grand Challenges. Additionally, the complexity of current and future hardware architectures requires the development of new computational methods and software technologies. Such advances have been at least as significant in advancing ER science as advances in computer hardware.

National Research and Education Network. ESnet is the principal DOE implementation of the NREN. ESnet ensures needed present and future connectivity among DOE researchers, experimental sites, and computational resources. ESnet leads the nation in adopting innovative network technologies that are useful to the DOE community and to the nation in general.

Information Infrastructure Technology and Applications. This element ensures that technologies needed to create collaborative research environments are pursued. Focus areas of this element include virtual reality, image understanding, intelligent system interfaces, databases, and object bases. Current efforts also focus on the tools needed to support energy demand and energy supply.

Basic Research and Human Resources. This element includes MICS efforts in applied mathematics, including tools needed for modeling, analysis, and numerical simulation of physical phenomena associated with energy systems. The basic research efforts in computer science are related to the ability to pursue scalable scientific problems via the development of programming models and tools, the management and visualization of scientific data, expanded libraries for parallel computers, and message passing techniques for distributed computers. The human resources component of this element is motivated to ensure needed future human resources in all MICS elements through investment in education and training.

Importance of Services Offered

The DOE MICS program supports some of the world's most advanced mathematicians, as well as computational and computer scientists. MICS researchers, like those of the other ER programs, are generally separated geographically not only from each other but also from the computational and experimental resources to which they require access. The MICS program, like other ER programs, relies on ESnet to provide access to collaborators and distributed resources, but ESnet also provides access to a state-of-the-art network for supporting research in distributed computing and advanced networking. The nature of these disciplines results in the vision, research, development, and implementation of an expanding number of increasingly sophisticated networking-based applications and tools.

A network that provides sufficient bandwidth, low latency, the provision of advanced networking services, and the ability to develop network applications and tools is critical to the success of the MICS program and principal investigators.

Through the last three decades, MICS researchers were some of the earliest developers and users of network services, such as email, remote access, and file transfer. More recently they are credited with the development and deployment of the Mbone and associated videoconferencing tools. MICS principal investigators are developing and utilizing advanced imaging and visualization techniques and other highly data-intensive and networking-demanding applications such as virtual reality. MICS researchers are investigating needed networking services, such as quality of service, to satisfy these applications. Multiple-site distributed computing is needed by and will be enabled by MICS collaborators to the benefit of the entire ER community and beyond.

Some of the examples of areas being pursued by MICS researchers are cited here. These examples do not represent the entire requirements of the MICS program, but are meant to provide a partial survey of the network-related requirements of MICS.

Data-Intensive Environments

Greatly improved connectivity among major instruments at the U.S. national laboratories, their users, and associated computational/storage facilities could dramatically change the nature of user collaboration. Distributed storage facilities and computing resources that could operate in a truly interconnected manner would enable real-time data analysis; tools that enable the aggregation of computing resources will positively affect both the methodology and time frame of data analysis. This potentially allows the real-time modification of the experiment itself by the principal investigator who is monitoring and analyzing the experiment in this new environment. Such developments would make a number of scientific instruments more effective, including accelerators, synchrotron light sources, micro-spectrometers, scanning confocal microscopes, electron microscopes, and similar data-intensive research apparatus.

Berkeley Lab's development of a Distributed Parallel Storage System (DPSS) has allowed investigation into the techniques and benefits of large amounts of distributed storage that can be treated in aggregation. This requires underlying network connectivity and services expected from the future offerings that likely need to be developed within the ER community. An early implementation of a related project called WALDO (wide-area, real-time, distributed large-object management system) supports distributed health-care related data; this application requires data streams at the T3 level. In parallel, a MICS-funded distributed data gathering and analysis effort in support of high energy and nuclear physics is already using an OC12 network connection.

Tele-Immersion and Collaborative Virtual Environments

Argonne National Laboratory, in conjunction with other DOE labs and university partners, is developing a prototype tele-immersion system called ManyWorlds. ManyWorlds is designed to enable multiple users to interact with each other and with simulations and databases as though they are in the same physical environment. Tele-immersion allows users to see, hear, and touch each other and the data in their environment. Tele-immersion relies on virtual reality display environments at each location and high-performance network connections between sites. These prototype systems require the management of nine types of data flows between sites. These data flows range from streaming audio and video to virtual-reality tracking data and interactive simulation updates. Each of these data flows have particular requirements for latency and bandwidth, determined by how they are used and perceived by the user. Advanced quality-of-service networking capabilities (i.e., the ability to reserve bandwidth and low latency channels in the network) ultimately need to be deployed for the widespread adoption of tele-immersion for DOE ER applications. High-end tele-immersive systems require OC12 (622 Mbps) and beyond, and very low latency transport.

Computational Steering

Oak Ridge National Laboratory and Sandia National Laboratory are experimenting with the creation of a distributed computing environment between their sites, allowing the DOE to address computationally intensive problems that are beyond the capabilities of the computing resource at any one site. Such an environment requires the ability to develop and implement an environment for program development, scheduling, data migration, and execution on distributed computing resources. Network demand for this area is difficult to predict—it depends heavily on the architecture of the problem being studied, but can easily reach OC48 (2.4 Gbps).

ORNL and SNL have demonstrated a prototype collaborative computational steering environment, CUMULVS, distributed across their three sites. Being able to explore "what if" questions interactively with large distributed computational experiments will accelerate the pace of scientific discovery in materials design and combustion. ORNL and SNL have begun work to integrate the EigenVR virtual reality interface with the CUMULVS collaborative steering environment, which could provide a more immersive, near-real-time environment for the computational steering of large distributed applications by teams of remote scientists. This further accelerates the process of scientific discovery by allowing teams of experts to brainstorm and interact with a simulation while it runs. Estimated bandwidth requirements approach OC48, 2.488 Gbps.

The Oak Ridge Complex (Y-12) is taking advantage of ESnet's ATM virtual circuit-switched fabric to carry encrypted traffic. Direct ATM AAL-5 constant bit rate (CBR) links are being supported between Oak Ridge and locations of interest to Defense programs. The links are successfully carrying high-bandwidth collaborative applications at DS3 and higher speeds. The capacity will be increased to OC3 during the coming year, with future connectivity needs scaling as high as OC24 (1.2 Gbps)

Distributed Computation

The MICS-funded efforts in the computational chemistry Grand Challenge project require the full breadth of ESnet services across the levels of the network protocol stack. For example, massively parallel computers will, for the first time, provide modeling capabilities for heavy-element compounds. The collaborators, located at ANL, PNNL, LBNL, Eloret Research Institute, Ohio State University, and Syracuse University will develop the needed computation techniques, including code development and tools, at the applications and systems level. ANL will serve as the prime High Performance Computing Research Center, while LBNL/NERSC will serve as the source of needed production cycles for this effort. The physical distribution among the principal investigators and their computation researchers requires development of network-based collaboration tools, distributed file systems, and an extensible computational chemistry environment. Additionally, the network must provide the connectivity for both the batch and interactive use of the massively parallel computers. The interconnection of the computing resources at ANL, PNNL, and LBNL/NERSC requires bandwidth on demand approaching OC3 for this collaboration.

In the future, this project will make full use of Mbone-based videoconferencing tools, including whiteboard and screen-capture capabilities. This requires development and deployment of the full suite of Mbone tools to be interoperable over popular computer platforms.

Needed file sharing requires sufficient bandwidth to ensure effective data transfer rates, accompanied by sufficient security mechanisms to ensure data privacy and protection. Analysis of such data will use modern visualization tools, putting high demand on the network in terms of both bandwidth and response.

The current situation places high demand on the future of ESnet and associated tools. Current network latency is insufficient to allow needed distributed collaborative development of needed tools, to the point where manual movement of source codes is more efficient than distributed code development. Demands for visualization and the requirement to remotely archive resulting data exceed the bandwidth currently provided by ESnet, forcing users to less-than-optimally select data for viewing or storage.

Computational Grids

A variety of next-generation energy research applications require the coupling of computers, storage systems, display devices, instruments, and people to form high-performance computational grids. These grids will feature integrated security, scheduling, communication, configuration, monitoring, and management services. These services will allow users to request and obtain computational resources regardless of location, and to construct networked applications that can meet stringent and dynamically changing requirements for end-to-end performance. The Globus project at ANL (in partnership with the Information Sciences Institute) is developing key grid technologies, deploying these technologies in prototype grid testbeds, and applying technologies and testbeds to challenging applications.

End-to-end performance guarantees for computational grids require integrated scheduling and instrumentation mechanisms on ESnet.

DOE2000

The MICS program, in conjunction with other DOE programs, is supporting a large-scale research and development effort to build tools and technologies for supporting the construction of collaboratories. This effort, known as DOE2000, is also supporting pilot projects in the use of collaboration technologies and an effort to design and build next-generation libraries for scientific computing.

The DOE2000 program has three main goals:

1. **Improved ability to solve DOE's complex scientific problems.** The scientific and technical problems faced by DOE now involve multiple complex sub-problems in several disciplines. For example, tracking radioactive waste plumes in groundwater requires contributions from fluid dynamics, chemistry, geology, and computational science. This has two implications. First, it is necessary to bring together scientists and engineers from multiple disciplines and organizations to collaborate in seeking solutions. Second, the nature of many problems precludes traditional experimentation because of the minute or enormous scales in time or space, the dangers to participants, or sheer unfeasibility.

2. **Increased R&D productivity and efficiency.** The panoply of problems to be solved, and the budgetary realities that frame possible approaches, require that greater productivity and efficiency in the problem-solving processes be achieved. While no one can schedule invention, DOE can do a better job of exploiting many centers of particular expertise and unique instrumentation or facilities. Tools that obviate geographical distance as a barrier to collaboration among people, and among people using remote instruments, will shrink travel time and expense. Mechanisms that allow teams to harness multiple, geographically separate high-performance computers will give more powerful simulation tools for problem-solving than ever before.

3. **Enhanced access to DOE resources by R&D partners.** External contact by people from academia, industry, and other governmental units to DOE's expertise and facilities enables us to form better collaborative teams. Those units in turn benefit from our collaborations in their R&D activities.

The DOE2000 program relies on ESnet to provide high-speed connectivity and advanced services in support of the collaboration technology development projects, collaboration pilot projects, and distributed computing libraries developed as part of the ACTS Toolkit.

The above examples, while discussed separately, need to be supported in aggregate.

The investigation and implementation of advanced networking services is also needed to provide the total offering of network related tools needed by such data-intensive applications. These include data cataloguing and multiple-site data distribution tools, easily configurable cache and processing systems, and security mechanisms.

Supercomputer Support

Energy Research is supported by production supercomputing at NERSC and research into high-performance computing at ANL, LANL, ORNL, and LBNL. The applications are run at the utilized center, but storage and certainly analyses may be accomplished at the site of the experiment, which may or may not be the site of the experimenter. This requires multiple channels of connectivity among the DOE sites, each of these channels requiring end-to-end quality of service (QoS). This type of QoS offering is not yet available but is recognized as a future user requirement. Additional research and development is needed to bring the QoS offering into prototyping and then production.

Connectivity Requirements

DOE MICS researchers require networking support to the following laboratories and universities (only universities that receive appreciable funding are listed here):

National Laboratories

Ames, ANL, BNL, JLab, LANL, LBNL, LLNL, ORNL, PNNL, SLAC.

Universities

University of Arizona; University of Auburn; California Institute of Technology; University of California, Berkeley; Colorado State University; University of Colorado; Cornell University; Florida State University; Harvard University; University of Illinois; Jackson State University; Johns Hopkins University; University of Maryland; Massachusetts Institute of Technology; University of Minnesota; New York University; Northwestern University; Princeton University; Purdue Research Foundation; University of South Carolina; State University of New York, Stony Brook; Texas Engineering Experiment Station; University of Texas; University of Washington; Western Washington University; University of Wisconsin.

International

Austria, Australia, Denmark, Japan, Norway, Poland, Sweden, United Kingdom.

Diesel Collaboratory Revolutionizes Design Process

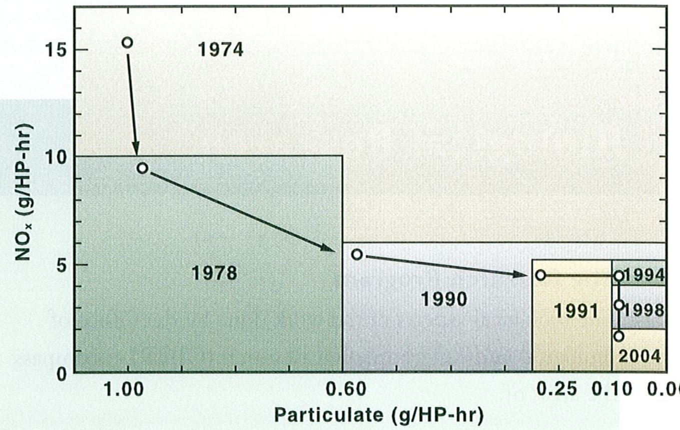

To meet increasingly strict regulatory goals for reducing nitrous oxide and particulate emissions, cleaner-burning diesel engines must be developed.

Since it was invented more than 100 years ago, the diesel engine has proved its reliability, durability, and relative efficiency. One of the engine's drawbacks, however, has been its exhaust emissions.

Today, ESnet is helping drive DOE's ongoing support of research to design and build more efficient and cleaner-burning diesel engines. By providing access to data and experimental facilities, allowing videoconferences and thereby fostering closer collaborations, ESnet is serving as a catalyst for researchers in the field.

The flagship effort is the DOE2000 Diesel Combustion Collaboratory (DCC), which aims to utilize computational models to streamline the design process with an enhanced rate of information exchange among scientists in diesel engine research.

ESnet is expediting this exchange of information between partners in the Department of Energy's Heavy Duty Diesel Combustion CRADA (Cooperative Research and Development Agreement). The network allows scientists at Sandia, Los Alamos, Lawrence Berkeley, and Lawrence Livermore national laboratories to more closely collaborate with their counterparts at the University of Wisconsin and three industrial leaders: Cummins Engine Co., Caterpillar Inc., and Detroit Diesel Corp. The Collaboratory is evaluating, testing, and implementing a set of tools designed to work across ESnet. These tools will allow researchers at these geographically distributed sites to:

- share graphical data from experiments or models using desktop computers
- discuss modeling strategies and quickly exchange model descriptions between groups
- archive experimental data, model data, and presentation materials in a Web-accessible electronic notebook
- produce a system to run combustion models on high-performance computers at widely separated locations
- build Web-based data processing tools to allow quicker analysis of experimental data and modeling results
- videoconference one-on-one collaborations and group meetings using desktop computers.

One of the major goals of the Collaboratory is to make modeling more efficient. By allowing direct access to the data archive, research partners not located at the national labs will be able to use the large computers at those sites more easily. From their own workstations, scientists will be able to discuss modeling strategies, exchange model descriptions between groups, and run modeling programs. Because much of the data generated by DCC will be proprietary, it will be necessary to implement security strategies into the data-sharing tools.

Videoconferencing capability is a key element in DCC. There are many instances when meetings of a team or subgroup would be productive, but time commitment and cost outweigh the usefulness of getting together. The proposed solution is a system that incorporates multi-site, workstation-based videoconferencing. Everyone in the group would have a microphone and camera at their workstation, enabling them to hear and see each of the other group members. Videoconferencing is a cost-effective way to foster lively interchange among scientists and engineers across the country, accelerating the pace of research and development. ESnet makes this progress possible.

Biological Research

The Research Program

The biological aspects of the work done by the Office of Biological and Environmental Research (OBER) encompass the areas of:

- human genome
- human subjects
- microbial genomes
- health effects
- structural biology
- bioinformatics infrastructure
- genome instrumentation
- medical applications.

While researchers in all these areas require network services for generic purposes such as email, the genome, structural biology, and bioinformatics areas have unique requirements. Indeed, they could not exist in their present forms without the Internet.

Every form of life 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 helical 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), recognition and marking of foreign molecules, forming gateways into cells, and providing structure and locomotion for cells. Knowing the structure of a protein is enormously helpful in understanding how it performs its function. Currently, precise and certain structural information can be gained only through lengthy and complicated experiments, but computational methods are being developed that can provide useful information in certain cases. Beyond knowing the structure of a protein, knowledge of its detailed interactions with other proteins and with other types of molecules is critically important for understanding cellular processes and for designing drugs that alter the protein's performance. This knowledge is also accessible through a combination of experimental and computational methods.

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 or other cells to incorporate specified DNA sequences and produce the resultant proteins on command. 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. The 1990s have seen large-scale use of DNA sequencing methods, including the determination of the complete genetic sequence (genome) of several organisms. Many more such determinations are in process, and the entire human genome is expected to be known by about 2004. In the area of structural biology, advanced synchrotron x-ray sources are coming on line that will aid and speed the experimental determination of protein structures.

Taken together, these developments have 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.

Knowledge of an organism's genome is virtually useless in isolation, however. To the eye, it is an incomprehensible sequence of the letters A, C, G, and T, anywhere from a few thousand (for viruses) to several billion letters long. The genome remains incomprehensible so as long as it is expressed on the printed page. The most convenient way to store and use genetic information is on a computer, preferably in some structured form, organized and annotated in a database. Additional information, obtained by experimental or computational methods, allows the interpretation of sequence information and makes it useful for a number of purposes. This additional information includes, for example, knowledge of:

- distinctions between functional and non-functional regions of DNA
- similarity of regions of DNA in an organism to other regions in the same organism
- similarity of regions of DNA in an organism to regions in other species
- identity of the protein or other product produced from a DNA region
- structure and function of the protein produced from a DNA region
- interactions of the protein produced from a DNA region with other proteins, small molecules, and with DNA itself
- observed mutations (variations in the DNA sequence) and their effects on an organism.

Information from several of these categories was used in designing HIV protease inhibitors, potent anti-AIDS drugs.

All these types of information, and many others, must be available to biological researchers for them to be able to do their jobs efficiently.

Connectivity Requirements

For historical and practical reasons, biological information is not kept in a single, monolithic database. Rather, it is distributed on computers at laboratories, scientific institutions, and corporations around the world, in various formats, according to the interests and inclinations of the researchers involved. Ubiquitous network connectivity is essential both to allow researchers to access individual pieces of information and to provide database maintainers the ability to cross-reference information in other databases. The new field of bioinformatics has arisen to service these needs, and bioinformatics infrastructure is one of OBER's research areas.

Compared to other areas of energy research, DNA sequencing (and to a lesser extent, structural biology) are small-scale science. A sequencing lab can be established for about $100,000, and it would be difficult to find a university or a national lab without one. While protein structure determination is more expensive, most universities, national labs, and biotechnology companies have such research programs. Growing interest in biotechnology as a scientific and commercial enterprise also contributes to the increase in sequencing and structure labs. Thus, widespread network connectivity is important to biological researchers, and high-capacity connections between the major database locations (and their mirror sites) would aid the integration of database functionality.

Biological research is also distributed internationally, and access to databases in foreign countries is essential to U.S. researchers. Current research takes place principally in the industrialized countries, and that is where the databases are kept. However, for reasons of both cost and interest, some developing countries are encouraging biotechnology research. A common focus is the genomes of local food crops and pests. These countries can be expected to play an increasing role in the near future and to contribute accordingly.

Capacity Requirements

While the network needs of biological research are met well by the current network backbone, in many locations connection to the backbone is a problem.

In planning for the future, recent trends may be of some help. Typically, the requirements of biological databases are increasing nearly exponentially, both in holdings and network connections served. This trend is illustrated for the "SWISS-PROT" database of protein sequences (exclusive of mirror sites) in Figures 1 and 2 (see page 35).

In the area of structural biology, new synchrotron x-ray sources, such as the Advanced Photon Source at ANL, are coming on line and will allow the collection of diffraction data at rates of multiple gigabytes per day. Currently, such experiments last at most a few days, and raw data is carried

New Tools Shed Light on Ancient Disease

In the hills above Berkeley and overlooking San Francisco's Golden Gate, Lawrence Berkeley National Laboratory's Advanced Light Source flashes like a beacon heralding a new era of scientific research.

This electron accelerator is the nation's brightest source of soft x-ray light. These tiny wavelengths of light are capable of imaging samples of materials as small as the wall of a red blood cell, just 1/5000th the width of a human hair. Among the many research projects under way at the ALS, biologists are using the light source to study one of the oldest and most persistent of all human diseases—malaria.

According to the World Health Organization, each year 300 to 500 million people living in the tropics and subtropics become infected with malaria, suffering burning fever and severe pain. Nearly 3 million—mostly children—die. Medical researchers have been unable to stamp out this scourge because the parasite's complex life cycle makes it an extremely tough opponent.

Cathie Magowan, a parasitologist in the Berkeley Lab Life Sciences Division, is using the x-ray microscopy beamline at the ALS to obtain never-before-seen views of the malarial parasite inside an intact red blood cell. Electron microscopes have been used before to obtain valuable information about the parasite inside a cell, but the cells had to be dehydrated and sliced into thin sections before they could be imaged. Consequently, there have been few detailed images of the malaria parasite within an intact cell. Using the ALS x-ray microscope, Magowan plans to record the entire life cycle of the parasite in red blood cells.

Already, Magowan and her colleagues in Berkeley and at Monash University in Australia have produced the clearest, most detailed images ever obtained of a malaria-infected red blood cell. As soon as the x-ray microscope produces the images, they are sent to a central database where they can be accessed by collaborators in Australia or other scientists around the world via the Internet.

ESnet is providing the technologies and services that allow flexible and economical online access to one-of-a-kind national research facilities like the ALS. This access results in better utilization of scientific facilities, more productive use of tight research budgets, greater diversity of ideas from multi-site collaborations, more efficient use of researchers' time, and the potential for even greater cooperation in addressing some of the most pressing scientific problems of our time.

X-ray microscopic images of the malaria parasite in a red blood cell are available online as soon as they are produced.

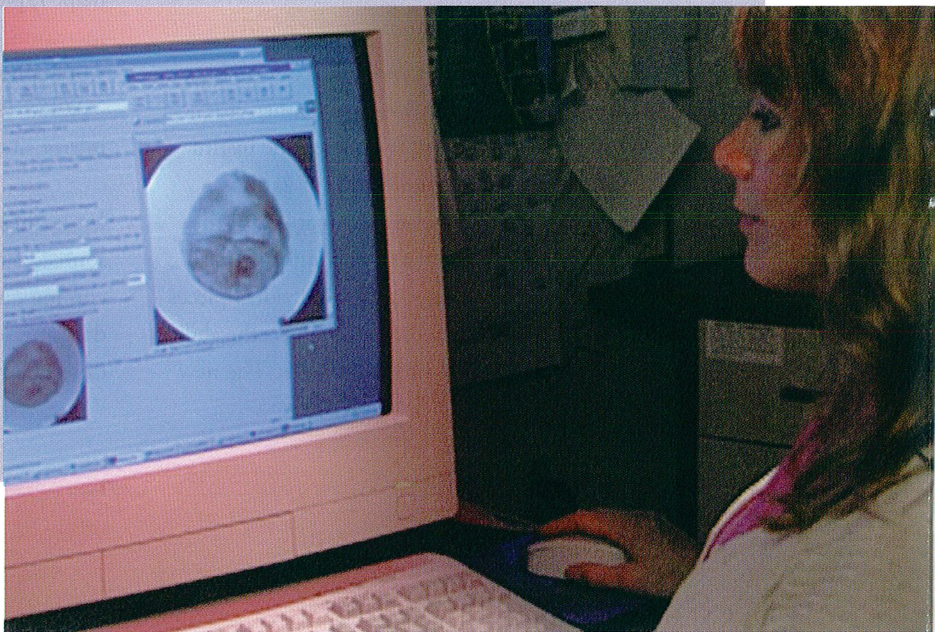

home on tape or other storage media. Reduced data sets, of considerably smaller (but growing) size, are frequently sent over the network, and their deposition in structural databases is now becoming common.

Collaboratory Requirements

While lacking the formal name collaboratory, the extraordinary cooperation of scientists in maintaining and using biological databases, including the use of shared tools for access and analysis, has many of the characteristics of a collaboratory. Its network requirements are described above.

Grand Challenge Requirements

The protein structure problem—the prediction of a protein's structure from its sequence—is one of the Grand Challenge problems. All methods of attacking this problem rely on massive computation, and some have no special need for network access. Other approaches rely on searching databases or requesting specialized services over the network (for example, protein secondary-structure-prediction servers). But because of the enormous mismatch between local disk access and network access speeds, the results are typically cached locally and have no special network needs.

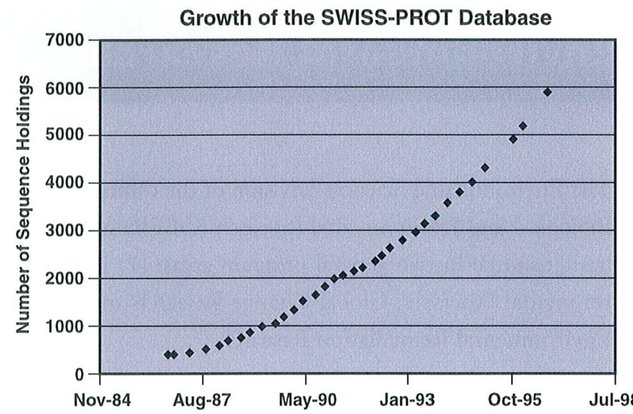

Figure 1

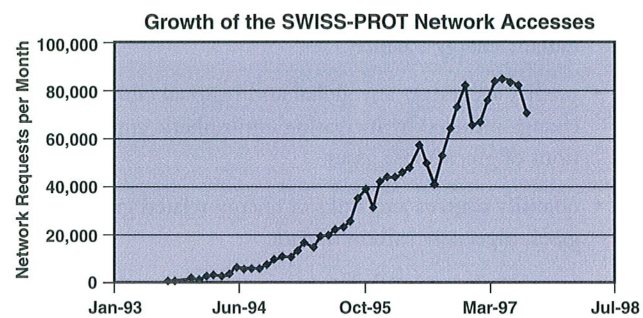

Figure 2

Environmental Research

The Environmental Sciences Division of the Office of Biological and Environmental Research (OBER) supports basic research in two general program areas: (1) Environmental Processes: Global Change Research and (2) Environmental Remediation Research.

Global Change Research

The Global Change activities include the process research and modeling efforts needed to:

- improve understanding of factors affecting the Earth's radiant-energy balance
- predict accurately any global and regional climate change induced by increasing atmospheric concentrations of greenhouse gases
- quantify sources and sinks of energy-related greenhouse gases, especially carbon dioxide
- improve the scientific basis for assessing the potential consequences of climatic changes, including the potential ecological, social, and economic implications of human-induced climatic changes caused by increases in greenhouse gases in the atmosphere and the benefits and costs of alternative response options.

While all of the Global Change research programs are served by ESnet, three with particularly demanding requirements are the Atmospheric Radiation Measurement Program, the Carbon Dioxide Information Analysis Center, and the CHAMMP Climate Modeling Program.

Atmosheric Radiation Measurement Program

The Atmospheric Radiation Measurement (ARM) Program is a multilaboratory, interagency program that was created in 1989 as part of DOE's effort to resolve scientific uncertainties about global climate change, with a specific focus on improving the performance of general circulation models (GCMs) used for climate research and prediction. These improved models will help scientists better understand the influences of human activities on the earth's climate.

In pursuit of its goal, the ARM Program establishes and operates field research sites, called Cloud and Radiation Testbeds (CARTs), in several climatically significant locales. Scientists collect and analyze data obtained over extended periods of time from large arrays of instruments to study the effects and interactions of sunlight, radiant energy, and clouds on temperatures, weather, and climate.

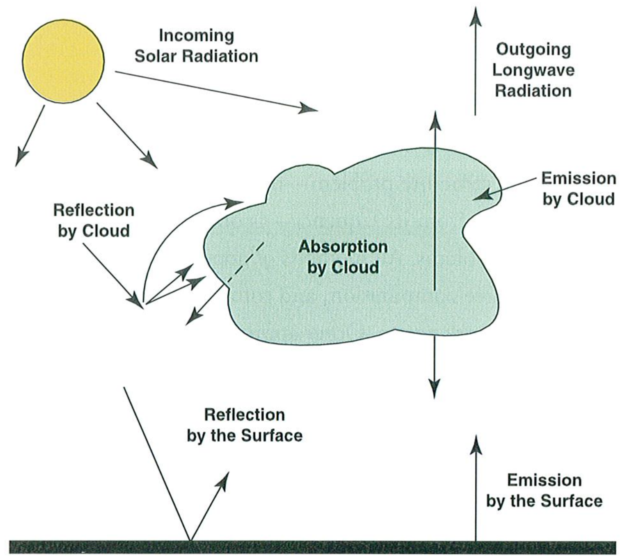

Understanding the Role of Clouds and Radiation

Three ARM CART locales are envisioned: one in the Southern Great Plains of the U.S., one in the Tropical Western Pacific, and one on Alaska's North Slope. Of these, the Southern Great Plains site and the first Tropical Western Pacific site, located off the coast of New Guinea, are operational. The first North Slope site, located in Barrow, Alaska, will deploy this year. ARM data flow collection, processing, and delivery is made possible by ESnet. ARM utilizes dedicated T1s for collection and communication to the Southern Great Plains and the North Slope. For communication to the Tropical Western Pacific site, an INMARSAT-B, 64 Kb satellite channel is utilized. As indicated in the following diagram, ARM's computing infrastructure is geographically dispersed. This year, the data flow through the ARM computing environment will grow to approximately 5 gigabytes per day; this number does not include data delivery to the ARM Science Team or general scientific community, both of which are variable depending upon research needs.

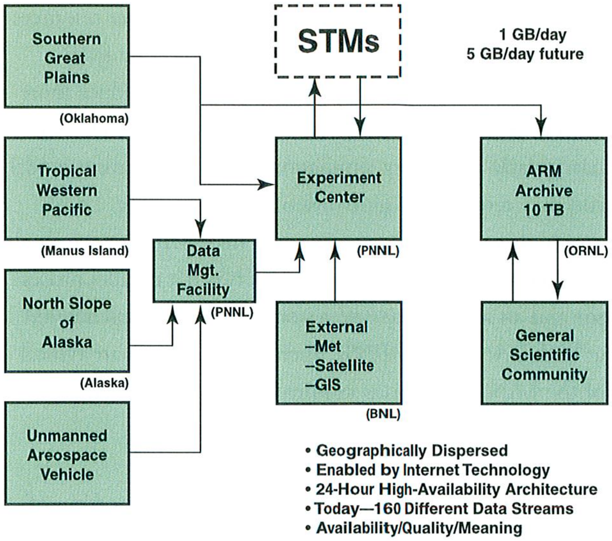

ARM Data Flow

The Data and Science Integration Team (DSIT) of ARM is composed of staff from several DOE national laboratories. The team provides both engineering and scientific support for data system development, implementation, operation, and maintenance for the computing, networking, science, and data processing needs of the ARM program.

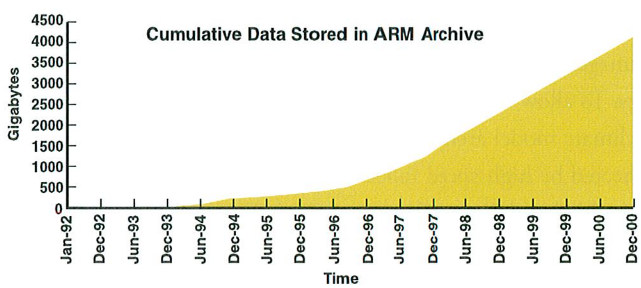

Projected Accumulation of ARM Data

Current data holdings of ARM exceed 1 terabyte and are expected to grow to 5 terabytes by the end of the century.

Translating science needs into data needs and delivering data streams of "known and reasonable" quality are fundamental principles of the ARM program. The measure of how well ARM addresses this vision is directly related to the satisfaction of the ARM Science Team through the scientific utility of its products. Further information on the ARM program is available on the World Wide Web at: http://www.arm.gov.

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, policy makers, 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 requesters from diverse settings in government and business.

Since its founding in 1982, CDIAC has responded to a hundred thousand 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 carbon-dioxide-related reports and produces a number of publications, including the newsletter CDIAC Communications, which, along with its global change databases, is distributed from CDIAC's World Wide Web site (http://cdiac.esd.ornl.gov) to thousands of users around the world.

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 13 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 ORNL Central Research Library, CDIAC also shares responsibility for operating one of the Regional Information Centers established by the International Geosphere-Biosphere Program.

The CHAMMP Climate Modeling Program

The Computer Hardware, Advanced Mathematics and Model Physics (CHAMMP) program is a DOE program to rapidly advance the science of decade- and longer-scale climate prediction. A major component of the CHAMMP program links the emerging technologies in high performance computing to the development of computationally efficient and numerically accurate climate prediction models. The program involves a joint laboratory-university effort to develop computational methods and simulation capabilities for future atmosphere and ocean general circulation models. These computer programs will form the core of advanced prediction models that can be used to study climate change.

Projects sponsored by CHAMMP are classed in two groups, Science Team and Development Team projects. The Development Team is implementing several state-of-the-art atmosphere and ocean general circulation models on high performance parallel computers. Recently the development teams have begun the process of coupling ocean, ice, and atmospheric components to study the utility of high-resolution eddy-resolving ocean models for the coupled climate system.

The CHAMMP World Wide Web site (http://www.epm.ornl.gov/chammp/) provides research results in the form of online technical documents and program information such as the CHAMMP newsletter.

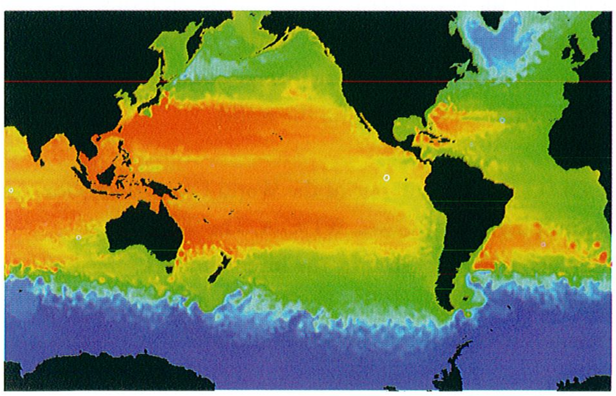

Sea surface temperature output from an eddy-resolving ocean model (the POP code) generated on a massively parallel computer at Los Alamos National Laboratory.

Networking Requirements For Global Change Research

DOE and the CHAMMP program support the nation's global climate research agenda by sponsoring the development of new models and providing dedicated computer resources for climate simulations ranging 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 gigabytes 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 to 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 National Center for Atmospheric Research in Boulder, the 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 staff at that site also disseminate the archived information to the general scientific community. However, most of the processing of experimental data occurs at PNNL. The data are analyzed at PNNL, 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 PNNL, 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 current and future operational field site increases (as it will when, for example, new imaging instruments come online), and as new field sites are brought online, 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.

Environmental Remediation Research

The Environmental Remediation Research portfolio is focused on developing an understanding of the fundamental physical, chemical, geological, and biological processes that must be marshaled for the development and advancement of new, effective, and efficient processes for the remediation and restoration of the nation's nuclear weapons production sites. A primary effort is a comprehensive research program in bioremediation that integrates the full range of fundamental scientific disciplines necessary to advance this emerging technology from one of hit-or-miss to one of sustained, cost-effective utility. Operation of the Environmental Molecular Sciences Laboratory, the only national scientific user facility dedicated specifically to DOE's environmental missions, is a key part of OBER's commitment to Environmental Remediation.

Natural and Accelerated Bioremediation Research

The mission of the Natural and Accelerated Bioremediation Research (NABIR) program is to provide the scientific understanding needed to harness natural processes and to develop methods to accelerate these processes for the bioremediation of contaminated soils, sediments, and groundwater at DOE facilities. DOE has a 50-year legacy of environmental problems resulting from the production of nuclear weapons. Among the most serious are widespread contamination of soils, sediments, and groundwater. The huge cost, long duration, and technical challenges associated with remediating DOE facilities present a significant opportunity for science to contribute cost-effective solutions.

While some of DOE's environmental remediation problems are shared by other federal agencies and the private sector, DOE faces a unique set of challenges associated with complex mixtures of contaminants, especially those mixtures that contain radioactive elements. In many cases, the fundamental scientific information needed to develop effective technologies is lacking. NABIR builds on a long history of fundamental research in the life and environmental sciences to develop a body of knowledge that will support environmental restoration of DOE's facilities.

Environmental Technology Partnerships

The Environmental Technology Partnerships (ETP) program fosters economic growth and competitiveness by encouraging the evolution from pollution control and waste management to pollution prevention and resource conservation. The program involves participation across the DOE Offices of Energy Research (ER) and Energy Efficiency and Renewable Energy (EE). ETP efforts in the Office of Biological and Environmental Research focus on fundamental bioprocessing and bioremediation research to better manage industrial waste, and on integrated assessment to identify high leverage payoffs for applications of environmental technology. The fundamental bioprocessing and bioremediation research includes research to better understand how microorganisms remove or transform contaminants. Integrated assessment research includes enhanced lifecycle analysis of the impacts and economics of new and/or enhanced technologies and processes, from raw materials to use and final disposal or degradation. This research also includes sophisticated, science-based risk characterization.

William R. Wiley Environmental Molecular Sciences Laboratory

The William R. Wiley Environmental Molecular Sciences Laboratory (EMSL), the DOE's newest national scientific user facility, is located at Pacific Northwest National Laboratory in Richland, Washington, not far from the Hanford site. The mission of the EMSL is:

- to provide advanced and unique resources to scientists engaged in research on critical problems in the environmental molecular sciences
- to educate young scientists in the molecular sciences to meet the demanding environmental challenges of the future.

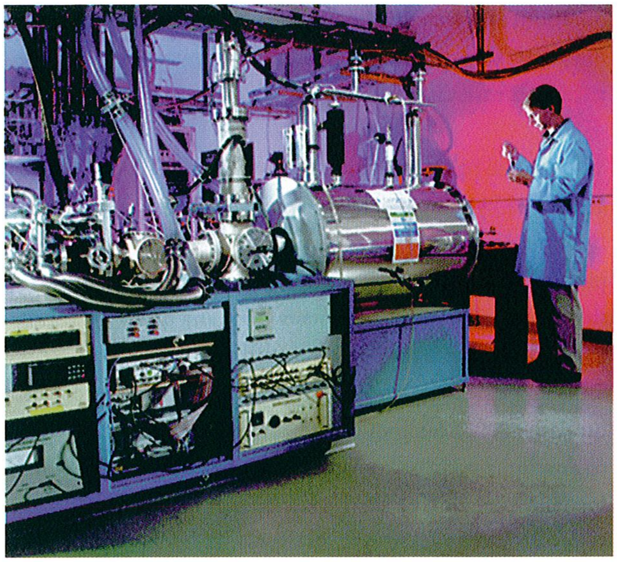

EMSL houses the world's first 11.5-tesla Fourier transform ion cyclotron resonance mass spectrometer.

As a research organization, the EMSL seeks:

- to attain an understanding of the physical, chemical, and biological processes needed to solve critical environmental problems
- to advance molecular science in support of the DOE's long-term environmental mission.

EMSL directly supports research for multiple offices of Energy Research, including OBER, OBES, and OCTR.

Through a groundbreaking approach to multidisciplinary environmental science, the EMSL provides its staff and users with a number of state-of-the-art research facilities. In many cases, the instrumentation in these facilities is unique (e.g., 11.5-tesla FTICR mass spectrometer); in other cases, the collection of research equipment provides a capability unsurpassed in other laboratories. Major facilities and capabilities include:

- The High Field Magnetic Resonance Facility supports studies of the molecular structure of enzymes, proteins, and DNA as they relate to bioremediation and cellular response effects. The facility contains a suite of state-of-the-art nuclear magnetic resonance (NMR) instruments, including a 750-MHz NMR spectrometer, two 600-MHz NMRs, and 500-MHz wide-bore NMRs with microimaging capabilities, as well as a suite of modern pulsed electronic paramagnetic resonance (EPR) instruments. An ultrahigh field NMR is also under development for the EMSL.
- The High Field Mass Spectrometry Facility focuses on biochemical applications of mass spectrometry, including combinatorial chemistry and sequencing applications. This facility contains a broad array of mass spectrometers, including the world's first 11.5-tesla Fourier transform ion cyclotron resonance (FTICR) mass spectrometer.
- The Molecular Science Computing Facility supports a wide range of environmental molecular research activities, from benchmark calculations on small molecules, to reliable calculations on large molecules and solids, to simulations of large biomolecules, as well as reactive chemical transport modeling. The MSCF features a 512-processor IBM RISC System/6000 Scalable POWERparallel (IBM SP) computer system with 67 gigabytes of memory and 2.9 terabytes of online disk, and a 20 terabyte EMASS hierarchical storage management system. A new generation of molecular modeling software has been developed to take full advantage of the advanced computing systems installed in the MSCF.

Other EMSL facilities support research with capabilities in

- nanostructural materials
- interfacial structures and compositions
- reactions at interfaces
- gas and liquid phase monitoring and detection.

Together, this combination of capabilities supports the full range of environmental molecular research, affording visiting researchers an opportunity to create their own unique user environments to address specific research problems. Staff members with specialized scientific expertise provide unique scientific capabilities in the EMSL. This expertise is key to the effective implementation of systems, tools, and facilities.

The Environmental Molecular Science Collaboratory leverages DOE2000 collaboratory technology research and EMSL's advanced computing infrastructure to enable geographically distributed teams of scientists to fully participate in experiment planning, data acquisition and analysis, report generation, and daily discussion, all from remote sites. The goal is to increase the efficiency of research, share EMSL expertise, data, and unique scientific instruments with remote colleagues, and allow a broader spectrum of complex research tasks to be undertaken.

Networking Requirements for EMSL

A hallmark of the integrated capabilities of the Environmental Molecular Sciences Laboratory is the partnership of theory and experiment. Thus, EMSL combines the networking requirements of a major computing facility with those of a large experimental user facility. Because of the wide range of research disciplines represented among EMSL's experimental capabilities, networking must meet the needs of a broad spectrum of research approaches in the molecular sciences. For EMSL's mission, computing and communications have become primary enabling factors, and the success of EMSL's environmental research is heavily dependent upon high-quality research interactions with scientists around the world, providing access to EMSL facilities and data, integrating expertise within and across domains, and supporting communication between DOE problem holders and researchers.

Connectivity: EMSL began full operations in September 1997, ramping up from a solid base of research programs and collaborations established during the construction phase. Now that the facility is complete, the capacity for research has been multiplied many-fold. EMSL can support over 200 resident staff, nearly 100 concurrent visiting scientists, and ongoing interactions with hundreds of remote scientists. Environmental molecular science researchers are distributed nationwide and worldwide.

This large and widely distributed user community creates requirements for high performance network connectivity to universities and federal labs throughout the nation and to strategic sites overseas. EMSL users come from programs throughout the Federal research arena, including DOE, NSF, NIH, etc., as well as industry. As the only DOE scientific user facility in the Pacific Northwest, EMSL also has a major role as a regional research resource. Thus, high performance connectivity to ESnet and (via peering relationships) to other national internets is paramount, and connectivity to regional networks in the Pacific Northwest is essential. Collectively, BER program requirements justify the consideration of PNNL as a primary ESnet backbone site.

Capacity: The ability of scientific user facilities to efficiently serve a wide community of collaborative and independent investigators has become a strong function of the capacity of network connections to those facilities. EMSL is in a rapid growth phase, and its network capacity needs will follow both the growth in users and advances in computing technologies. For example, EMSL's Molecular Science Computing Facility features not only one of the most powerful computers in the world, but also a new generation of applications software, including the NWchem high performance computational chemistry codes and the Extensible Computational Chemistry Environment (ECCE'), an integrated problem-solving environment for computational chemistry.

The Molecular Science Computing Facility contains a 512-processor IBM RISC System/6000 Scalable POWERparallel computer system, supporting a wide range of environmental molecular research.

MSCF projects are typically large team projects, involving several institutions. The majority of MSCF users will be remote, depending on rapid network access to MSCF systems and data, and employing distributed computing approaches to running problem-solving environments. On the experimental side, EMSL's High Field Magnetic Resonance and Mass Spectroscopy Facilities both feature instruments which create large datasets. A significant fraction of the experiments with these instruments will include remote operation. For example, exciting new opportunities in mass spectroscopy involve real-time data analysis, where a subset of the dynamically trapped ions is selected for the next phase of the experiment by analyzing data while the experiment is going on. Data transport and visualization for these kinds of experiments is intensive and time critical.

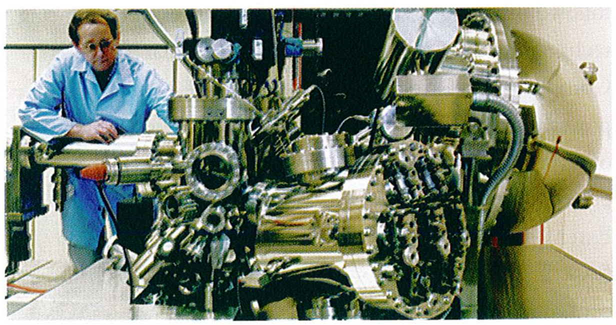

A molecular beam epitaxy system at EMSL supports molecular-level studies of interfacial and intraphasial chemical reactions.

Access to EMSL's computers, instruments, and databases, combined with the collaboratory requirements below and other BER programs at PNNL, make a solid case for 150 Mbps network service to PNNL within a year and 600 Mbps by the turn of the century.

Collaboratory: Research collaborations are a hallmark of environmental science, and electronic collaboration technologies are opening up many new opportunities for assembling the teams and facilities needed to address the fundamental research issues in environmental remediation. The Environmental Molecular Sciences Collaboratory is one of several DOE projects to demonstrate the value of this approach; however, it is clear that present ESnet capabilities need to be scaled up rapidly to meet the developing collaboratory demand. Shared visualization, collaborative instrument control, shared screens, and videoconferencing are integral parts of collaboratory environments at EMSL. With hundreds of simultaneous research projects going on involving EMSL, and a substantial fraction of the offsite investigators using collaboratory tools regularly, the aggregate requirement for network resources is expected to be quite large. We must also keep in mind that end-to-end performance is the crucial metric, where EMSL's remote collaborators come from many institutions. Therefore, it is important that ESnet services connect seamlessly to and interoperate readily with other networks.

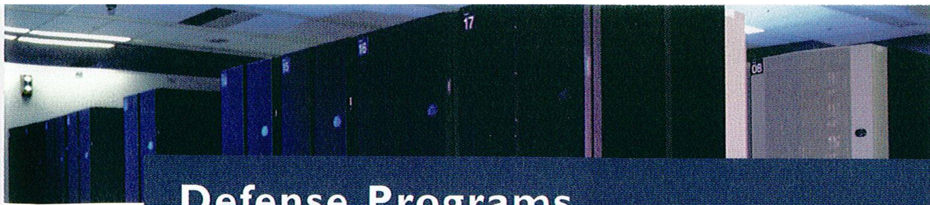

Defense Programs

The mission of Defense Programs (DP) is to ensure the safety, reliability, and performance of nuclear weapons without underground nuclear testing. DP provides an infrastructure and the intellectual capability to maintain the nuclear weapons stockpile, including replacing limited-life components and assuring an adequate supply of tritium. DP also provides the ability to reconstitute underground nuclear testing and nuclear weapons production capabilities as required to meet future national security requirements.

DP has initiated efforts to develop technologies in support of its national security mission under requirements of the Comprehensive Test Ban Treaty (CTBT). This is to ensure the performance, safety, and reliability of the nuclear stockpile, while adhering to the "zero-yield" provisions of the CTBT. DP is bringing about a new era of Science-Based Stockpile Stewardship (SBSS) based on advanced modeling and simulation methods, complemented by aboveground experiments (AGEX) and validated against historical test data.

Accelerated Strategic Computing Initiative

Future weapon assessments must rely on the judgments of technical personnel increasingly removed in time and experience from nuclear testing. This will require continual enhancement of computational technology for simulation and modeling. In particular, the Accelerated Strategic Computing Initiative (ASCI) was established as a critical element of stockpile stewardship to shift promptly from test-based to science-based assessment methods, and will create advanced computational capabilities essential for success. The data produced by new ASCI 3D fuller-physics applications will vastly exceed today's in both quantity and complexity, and will require innovative approaches to data analysis and data assimilation tasks.

The ASCI program is executed at Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratories and with supporting research at five University Alliance Centers: California Institute of Technology, Stanford University, University of Utah, University of Illinois, and University of Chicago. Manufacturing factilites are located at Oak Ridge National Laboratory, Pantex, Allied Signal, and Savannah River.

ASCI employs five strategies to achieve its objectives:

1. Create seamless management: One Program—Three Laboratories.
2. Focus on advanced applications development.
3. Focus on the high end of computing.
4. Create problem-solving environments.
5. Encourage strategic alliances and collaborations.

Create Seamless Management: One Program—Three Laboratories

The problems that ASCI will solve for the Stockpile Stewardship and Management Program span the activities and responsibilities of the three Defense Programs laboratories. Cooperation is essential to the stockpile Vision 2010. The DP laboratories participate in the ASCI program as partners.

In accordance with this philosophy, the ASCI program was developed with full participation from representatives of the laboratories. There will be unprecedented cooperation among the three DP laboratories in the future as well. The ASCI program is implemented by project leaders at each of the laboratories, guided by the Office of Strategic Computing and Simulation under the Assistant Secretary for Defense Programs. Finally, the ASCI research projects share code development, computing, storage, and communications resources across laboratory boundaries in joint development efforts.

Focus on Advanced Applications Development

ASCI is developing the high-performance software applications needed to implement virtual testing and prototyping. Tightly integrated code teams develop these codes. Code teams are large interdisciplinary work groups whose objective is to produce coherent software packages for efficient predictive simulations. ASCI application codes are the key to the ability to reach the Stockpile Stewardship and Management Program objectives outlined for 2010 without nuclear testing and at an affordable cost. ASCI will provide simulations embodying all the physics needed to predict the safety, reliability, performance, and manufacturability of weapons systems.

It is a formidable challenge to replace the empirical factors and adjustable parameters used in current calculations with predictive physical models. This challenge will produce large, complex applications that will drive the scale of computing machinery and the infrastructure. However, increased capability of these elements alone is insufficient. Much of the increased computational capability must come from improvements in the applications codes themselves.

The applications code development effort for ASCI focuses on applications in weapons physics, engineering, and manufacturing science. The combination of fundamental physics models, advanced 3D numerical methods, and high-end scalable computer systems requires substantially more complex applications codes than any developed to date.

These applications integrate 3D capability, finer spatial resolution, and more refined physics. This obviates the need for verification of results through underground nuclear tests and prototype performance validation tests. The nuclear performance applications will be verified and validated using data from a range of non-nuclear test facilities and experimental results from historical underground tests. These applications will be designed for maximum performance on the new scalable computer systems.

Focus on the High End of Computing

More powerful computers are needed for virtual testing and prototyping applications. ASCI will stimulate the U.S. computing industry to develop high-performance computers with speeds and memory capacities thousands of times greater than currently available models and tens to hundreds of times greater than future computers likely to result from current development trends. ASCI partners with various U.S. computer manufacturers to accelerate the development of larger, faster computer systems and software that will run DP applications more efficiently.

Create Problem-Solving Environments

ASCI requires an unprecedented code development effort at the DP laboratories. The problem-solving environment supports the applications code teams by providing a robust applications development environment with appropriate software engineering tools. To make the applications on computers usable at desktops throughout the DP laboratory complex, ASCI helps develop a computational infrastructure, including local-area networks, wide-area networks, advanced storage facilities, and software development and data visualization tools.

Encourage Strategic Alliances and Collaborations

ASCI requires the technical skills of the best scientists and engineers working in academia, industry, and with other government agencies. The need to develop an unprecedented level of simulation capabilities requires strategic alliances. These alliances are required to understand and support the development and credible demonstration of this simulation capability. Wherever possible, ASCI collaborates with organizations conducting research of interest. ASCI also works with existing organizations to develop and apply standards. Finally, ASCI plans to initiate exchange programs to bring top researchers directly into the project while allowing laboratory personnel to expand their experience base in external projects.

Service Requirements

The primary service drivers are to support remote usage of teraflops-scale computer systems, including the remote visualization of terabyte-scale datasets, and to support a geographically distributed information enterprise. These systems will be used in both classified and unclassified environments and by users located at the laboratories, University Alliance Centers, and plants. The ASCI Problem Solving Environment strategy and the Distributed and Distance Computing strategy are creating a distributed environment at the laboratories that incorporates distributed file systems, authentication, and resource management. Developing this distributed environment to meet the stockpile stewardship time constraints requires a partnership with the network service provider so that advanced technologies can be made available sooner than under a normal commercial evolution.

Connectivity Requirements

DP requires reliable high-speed networks that interconnect the three laboratories, five University Alliance Centers, and four manufacturing plants with the Internet. In addition, secure, reliable, high-speed networks that interconnect the laboratories and plants are also required.

Secure Virtual Network Supports Nuclear Stockpile Stewardship

In today's world, concerns about security seem to go hand-in-hand with using a computer connected to a network. And when the network is being used to conduct classified national security research, those concerns take on truly global implications.

Recognizing this, ESnet now supports an important new concept called "virtual private network" which enables an independent network to run transparently over an existing host network. This capability allows different virtual networks to operate on the same infrastructure—each one running independently from the other, each one with its own set of unique needs and technical requirements.

SecureNet supports classified research with protected access to the ASCI "Blue-Pacific" IBM supercomputer at Lawrence Livermore National Laboratory as well as ASCI systems at Los Alamos and Sandia.

One such "virtual private network" is SecureNet, a classifed network that links classified research supercomputers at Lawrence Livermore, Los Alamos, and Sandia National Laboratories and requires secure data transmissions. SecureNet's mission is to support the nation's Stockpile Stewardship Program and ASCI, the U.S. Department of Energy's Accelerated Strategic Computing Initiative. SecureNet transmits encrypted data between the three labs' supercomputers, DOE headquarters in Washington, D.C., and Defense Program production facilities—Allied Signal in Kansas City, Pantex in Amarillo, Texas, Oak Ridge National Laboratory, and the Savannah River Site.

By using the ESnet infrastructure and broad bandwidth capacity, researchers are able to create and run billion-node-mesh codes that maintain the safety, security, and reliability of the United States' nuclear weapons stockpile. Even if they aren't on-site, researchers have access to any one of the three ASCI supercomputers: the Intel system at Sandia, the IBM system at Lawrence Livermore, or the Cray system at Los Alamos. Complex codes that used to take several days to run can now be done in a fraction of that time. And because the network is independent and secure, information can be shared more quickly, but just as securely as in the past.

"It was very difficult in the past to collaborate with another laboratory, because just to pass a note on a classified project, you'd have to seal it in a package and send it by classified mail or find someone to hand carry it," says Randy Christensen, Deputy Program Manager for ASCI at LLNL. "The time span would be days between reiterations. You can't do collaboration in that environment. But with SecureNet, you can send it by email and have a response within minutes. It allows a really interactive, collaborative character between the projects that we didn't have before."

Also in development is a SecureNet Distributed Computing Environment, in which the three teraflops computing systems would work in parallel to handle even bigger calculations more quickly.

Capacity Requirements

The ASCI platform goals include a 10-teraflops system in 2000, a 30-teraflops system in 2001, and a 100-teraflops system in 2004. These systems will be used primarily for multihour simulations that will generate hundreds of terabytes of data. Since there will be one of each of these systems, they will be unique resources and will thus be accessed remotely. Providing sufficient bandwidth and minimizing the latency to support the remote usage of these systems will require new approaches to remote visualization.

The computing environment for FY99 is projected to include a 1.8-teraflops system and two 3-teraflops systems. The existing network connections to the laboratories are at OC3 rates and are shared between classifed and open networks. End-to-end performance, in terms of both bandwidth and latency, are the metrics that users see. The goal for FY99-FY00 is to sustain a rate of 100 gigabytes per hour and to achieve a latency of less than 1 millisecond per 100 miles between sites. Achieving these goals will require close collaboration between the ESnet team and the laboratories.

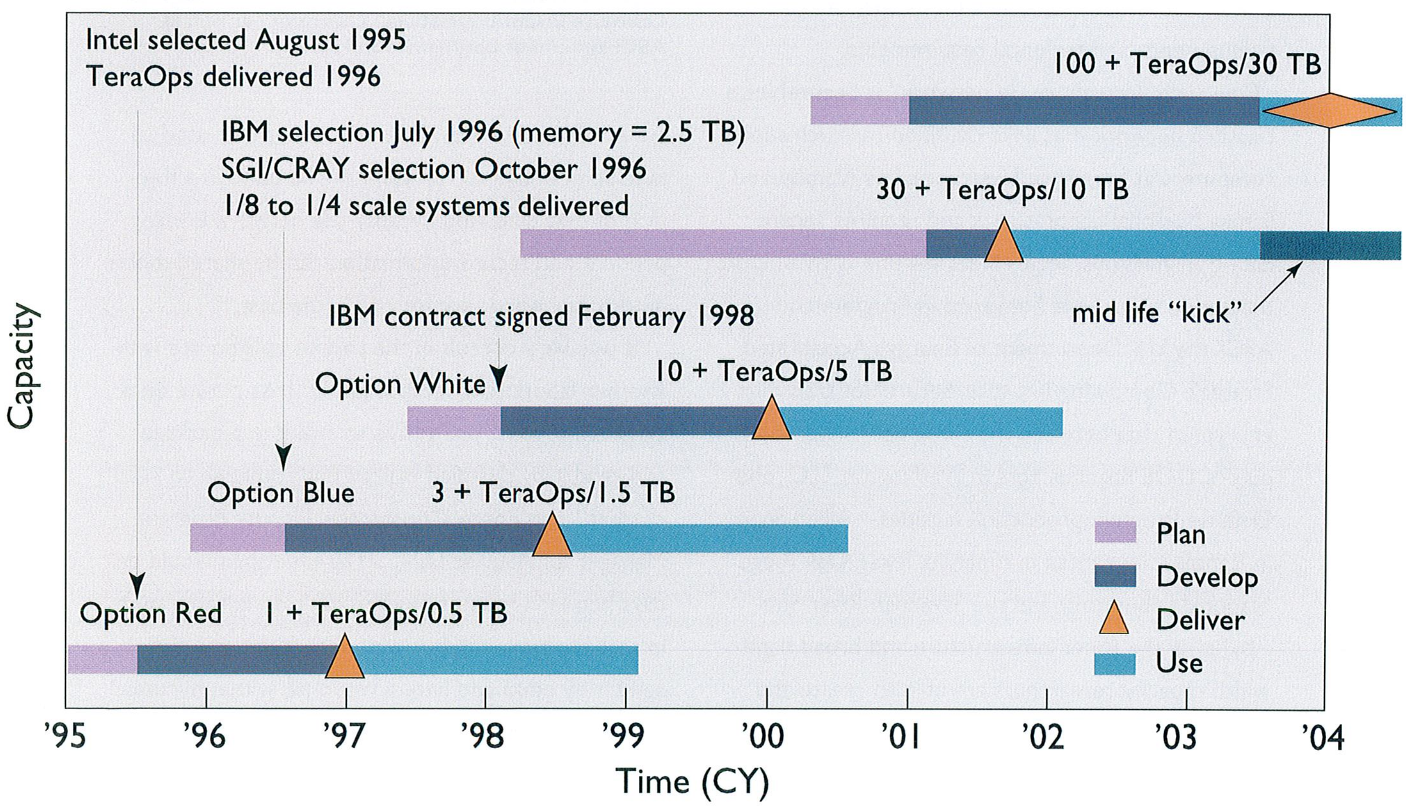

ESnet Organization and Processes

A successful service-oriented community must first establish user requirements and then effectively satisfy those requirements within the operating, funding, and administrative constraints. For ESnet, these processes involve one of the most widespread collaborative efforts within the Department of Energy. Assessments of current and future user requirements must be obtained from DOE-supported principal investigators located at national laboratories and universities across the U.S. and at other research sites throughout the world. The effort to fulfill these requirements is concentrated in the ESnet Project, located at Lawrence Berkeley National Laboratory, although the efforts of network specialists at all major DOE sites are also required.

The breadth of services provided by ESnet—from basic bandwidth to directory services to email 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.

The Mathematical, Information, and Computational Sciences Division (MICS)

ESnet is sponsored by the Mathematical, Information, and Computational Sciences Division (MICS) within the Office of Computational and Technology Research (OCTR) of DOE's Office of Energy Research. MICS is therefore responsible for funding the network and overseeing its management. MICS'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, MICS is responsible for managing ER supercomputer facilities and ESnet.

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 Offices of Energy Efficiency (EE), Nuclear Energy (NE), Energy Information Agency (EIA), and Human Resources (HR) 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 MICS, 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 project management personnel.

Members of the ESSC represent specific DOE program offices and are appointed by those offices, 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. MICS nominates an ESSC member to serve as committee chairperson, and the selection becomes final upon approval by a vote of the entire ESSC. These processes have been developed in the formative stages of the ESSC and are documented in the minutes.

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. This documentation defines the role of MICS 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 ESSC 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 foreign connectivity and bandwidth
- the need for DOE-wide coordination of distributed computing services
- the future effects of other networks on the ER community
- the relationship of ESnet to national initiatives
- the effects of related Federal activities, such as the Federal Networking Advisory Committee and the Federal Internetworking Requirements Panel.

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

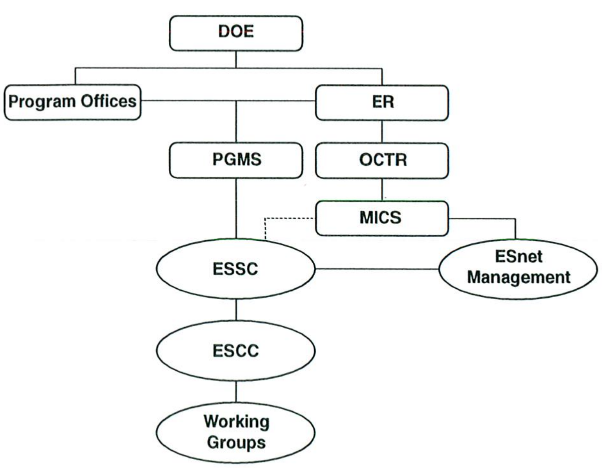

Relationships among DOE and networking entities: This figure represents a hierarchial view of the many entities involved with networking within DOE.The rectangles represent organizational components of DOE, and the ellipses signify structures within the ESnet community.

- 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 (ESCC) 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 MICS 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 ESSC 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 ESSC approval.

ESnet and the Larger Internet Community

ESnet has been and continues to be a partner with the larger Internet community. ESnet has historically been a leader within that community, having early explored new technologies such as dual protocol routers and ATM. By providing these successful pioneering implementations, ESnet helps advance the state of technology within the wider community.

The guiding principle behind ESnet's participation in the Internet community is one of positive participation wherever reasonable and feasible. Some limitations are, however, dictated by mission requirements, financial constraints, and proper regard for the security of computer resources served by ESnet.

ESnet personnel participate in the Internet community forums, task forces, working groups, committees, and subcommittees wherein standards, protocols, and acceptable practices are developed. Examples of these are:

- Internet Engineering Task Force (IETF)
- International Committee on Future Accelerators-Network Task Force (ICFA-NTF)
- Internet Statistics and Measurements Analysis (ISMA)
- North American Network Operators Group (NANOG)
- Joint Engineering Team (JET)
- ATM Forum.

Technical issues on which the ESnet community is providing expertise and leadership are:

- Internet security
- congestion analysis
- public key infrastructure
- IPv6
- operations management
- directory services
- privacy
- network research
- quality of service (QoS)
- applications.

There are many management and joint engineering opportunities with other parts of DOE, other government agencies, and the international research community that result in mutually beneficial arrangements. These are detailed in the following sections. ESnet is widely recognized for the technical competency of its staff and willingness to look for solutions to technical, financial, and political difficulties.

Memoranda of Understanding

Through the budget process, ESnet funding is appropriated to support Energy Research mission requirements and not for support of other major departmental programs or organizations unless otherwise provided for through Memoranda of Understanding (MOU). In recognizing certain goals of the National Information Infrastructure, such as improving the availability of and access to government information, other benefits to the DOE can be realized through the coordinated and cooperative use of the network infrastructure and its capabilities by other departmental programs and organizations. Such use is encouraged to manage overlap and to complement other communications requirements.

The accompanying descriptions give the mission of other DOE organizations and the role that ESnet plays for that organization as formalized by an MOU.

Energy Information Agency

Mission: Energy Information was established to develop and support a comprehensive, unified, and credible energy information and analysis program within the Executive Branch. Its mission is to inform DOE, Congress, the Executive Branch, and the public regarding the U.S. energy situation by administering a central and comprehensive program for collecting, interpreting, validating, analyzing, and disseminating energy information. The dissemination of information is one of Energy Information's major functions. It carries out a comprehensive dissemination program of energy statistical and forecast information. Both public and private sectors of the economy rely on Energy Information for all types of energy information.

Energy Information's Use of ESnet: Network capabilities aid Energy Information in accomplishing the following departmental information management objectives:

- To optimize availability, use, and benefit of DOE's energy information to all appropriate customers, including academia, industry, and the general public.
- To enable exchange of energy information among headquarters, field, and contractors.
- To maximize the sharing of energy information among DOE elements, contractors, other government agencies, and international partners to help ensure maximum program advancement and optimal use of resources.
- To encourage sharing of information resources within DOE, while minimizing duplication of information available elsewhere.
- To facilitate access to electronic mail and bulletin board services for the dissemination of information among DOE, Congress, the Executive Branch, and the general public.

Human Resources and Administration

Mission: Information Management has the primary responsibility within DOE for plans, policies, and oversight with regard to the management and use of information and telecommunications, as well as unclassified computer security and records management. This office also provides information technology and information management support to DOE headquarters organizations and field elements, as required. This support includes requirements analysis, acquisition, implementation, and maintenance support covering a broad spectrum of both computing- and communications-based hardware and software, as well as policy advice on the various functional areas.

Information Management's Use of ESnet: The Deputy Assistant Secretary for Information Management has sponsored task forces and studies to determine how to improve the DOE capabilities for data interchange and interconnectivity. Several recommendations were proposed, with the two most important being:

1. Long-term, to continuously implement community accepted capabilities as appropriate.
2. Near-term, to expand the base of data communications interconnectivity through the shared use of existing resources such as ESnet.

Energy Efficiency and Renewable Energy

Mission: Energy Efficiency is responsible for research, development, and deployment of energy efficiency and renewable energy technologies. Energy Efficiency executes its mission through work at DOE's national laboratories, private sector industrial firms, Federal agencies, and universities. A major feature of Energy Efficiency programs is the large number of cooperative development and deployment partnerships with industry and with state and local governments. These partnerships are facilitated by laboratories and by the Regional Support Offices that Energy Efficiency maintains in the ten federal regions. Energy Efficiency attempts to maintain information flow among these participants to enhance management processes and the provision of information to, and feedback from, major stakeholders in Energy Efficiency's technologies.

Energy Efficiency's Use of ESnet: Network capabilities will aid Energy Efficiency in accomplishing the following departmental information management objectives:

- To facilitate the creation, use, and value of scientific, technical, and educational information that furthers DOE objectives, such as promoting a more productive and competitive economy, improved environmental quality, a secure national defense, and enhanced education and scientific research.
- To optimize availability, use, and benefit of DOE's scientific, technical, and educational information to all appropriate customers, including academia, industry, and the general public.
- To enable exchange of technical data and information among headquarters, field, and contractors.
- To maximize the sharing of scientific, technical, and educational information among DOE elements, contractors, other government agencies, and international partners to help ensure maximum program advancement and optimal use of resources.
- To encourage sharing of information resources within DOE while minimizing duplication of information available elsewhere.
- To facilitate access to electronic mail and bulletin board services for the exchange of information among national laboratories and program participants in a variety of scientific disciplines.

Nuclear Energy, Science and Technology

Mission: The mission of the Nuclear Energy Program (NE) is to provide technical leadership to address critical domestic and international nuclear issues. NE contributes to energy supply diversity and advances U.S. competitiveness and security by providing nuclear products and services that meet the needs of the nation and the world community in an environmentally sound and economic manner. NE encourages public involvement in its programs and provides information to increase public knowledge.

NE Use of ESnet: NE has responsibility for ensuring availability of isotopes for health care, research, and industry; supplying radioisotope thermoelectric generators (RTGs) for exploration of deep space; developing advanced nuclear reactor and nuclear fuel cycle concepts; ensuring the safety of DOE reactors; and enhancing the safety of Soviet-designed reactors through the International Nuclear Safety Program (see below).

Because ESnet connects to virtually the whole of DOE, it is the ideal network to directly support the activities of the Nuclear Energy Program. Most of the program activities involve multiple national labs and DOE headquarters. To support advanced applications such as videoconferencing, World Wide Web-based information access, and other real-time applications, it is essential that NE utilize a network that poses the least variation in latency—which can only be ESnet, unless NE were to build its own network. ESnet's international connectivity is ideally suited for coordinating NE's efforts with the International Atomic Energy Agency and the European Community. There are alternatives to using ESnet, but they would involve unnecessary duplication of effort and expenditure.

International Nuclear Safety Program (INSP)

Mission: The International Nuclear Safety Program (INSP) is designed to reduce the national security and environmental threats posed by the operation of unsafe and aging nuclear power reactors around the world. Today there are 67 Soviet-designed reactors (SDRs) operating in Russia, Ukraine, Armenia, Bulgaria, the Czech Republic, Hungary, Lithuania, and Slovakia. Many of these reactors have deficiencies in safety equipment, training, safety procedures, and oversight. INSP seeks to improve the physical condition of the nuclear power plants, upgrade the training of plant operators, and promote modern safety technologies and methods. Working partnerships are established in host countries with regulatory and nuclear power plant personnel, government agencies, scientific institutes, and private contractors. INSP program goals include:

- Improving the safety and reducing the risks of operating SDRs.
- Assisting Russians in stopping production of weapons-grade plutonium and minimizing proliferation risks of weapons-useable nuclear materials.
- Reducing the national security and environmental threats posed by the destroyed reactor and the two remaining operational reactors at Chornobyl.

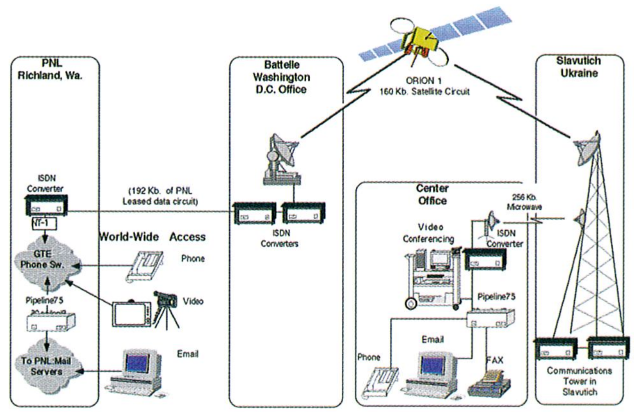

Chornobyl Center Satellite Communications System

INSP Use of Esnet: INSP has begun to use computer networking as a tool to help in the execution of its program and projects. The first networked computer systems were installed in Slavutych, Ukraine in August 1996 to support the Chornobyl Center for Nuclear Safety and Radio Ecology. This system supports network, videoconferencing, and telephone from Slavutych to PNNL in Richland, Washington. From Richland, the Slavutych system can be accessed anywhere in the world via ESnet. The Chornobyl Center is expected to be a focal point for the international initiative to construct a new shelter for the destroyed reactor and to decontaminate and decommission the remaining operational reactors.

Several major INSP projects over the next few years will rely on ESnet services:

- Supporting construction of a replacement shelter for the destroyed reactor at Chornobyl.
- Supporting the decommissioning of outdated reactors in Russia and other former Soviet Union countries.
- Providing training for nuclear reactor personnel in host countries, and tools such as simulators and other resources to improve safety at reactor sites.
- Replacing nuclear cores and converting reactors in Russia to domestic production.
- Providing database systems for sharing operational and maintenance information among all RBMK-design reactor sites in Russia, Ukraine, and Lithuania.

ESnet and Other U.S. Agencies

Scientific research within the U.S. is an activity funded by several government agencies. DOE-funded researchers participate in collaborations involving researchers funded by these agencies. This situation is fertile ground for cost-saving cooperation among agencies. Rather than having multiple lines to an institution of interest to multiple agencies, it clearly is desirable to have one agency provide the connectivity, let the others use it, and on a quid-pro-quo basis take advantage of the others' connectivity to a different institution of interest.

ESnet has an admirable history of participating in cost-sharing arrangements with NASA and NSF. The current ATM contract was negotiated jointly with NASA.

ESnet uses both the West and the East Federal Interchange Points (FIXs) to connect to the other agencies.

ESnet International Connectivity

Scientific research is a worldwide activity, and DOE-funded researchers participate in collaborations involving dozens of countries. Given international travel's expense and disruption to routine, the more effective the international connectivity, the greater is the benefit to the research program. The domestic connectivity and success of ESnet has made ESnet the network of choice for foreign researchers wishing to reach their American colleagues.

However, achieving the same degree of international connectivity as domestic is a goal still being sought. Several unique problems affect international connectivity:

- Providing bandwidth across the oceans is technologically challenging and expensive.
- Tariffs and infrastructure differ among the partnering nations.
- The policies of some partnering institutions limit the possible network configurations.
- Within DOE, the different research programs have varying connectivity requirements country by country, and this makes the process of prioritizing requirements within the user community cumbersome.
- Traffic from one country to another may pass through a third country, e.g., from Germany to Japan via America, or from America to China via Japan.

Currently, ESnet has international connectivity as follows:

Location	Bandwidth
Europe via DANTE	T3
CERN	E1
Italy	T1
Germany	T1
Russia	<T1
Japan	T1

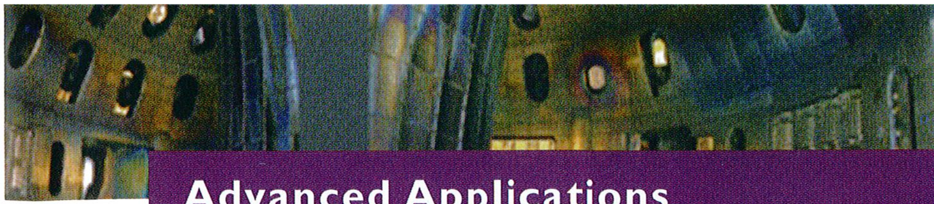

Advanced Applications

This section makes an initial identification of advanced networking technologies that are expected to underlie the future success of DOE science. This information is provided in an attempt to influence future networking services as well as associated research and development. Seven application areas, with appropriate development, promise to serve as a foundation for expected expansion and extension of DOE programmatic research capabilities. Five cross-cutting network technologies will enable and/or accelerate further scientific application advancements. This section closes with a short discussion of network research and development and the benefits of a collaborative approach.

The success of ESnet is ultimately determined by the advancements in DOE science that the network facilitates. ESnet providers are to be commended on their success to date within budget constraints and given agreed prioritization.

Past experience has demonstrated that networking advances are both application driven and technology driven. ESnet must continue to be responsive to both. The DOE community has applications that the next generation of network services must support. Synergistically, network-based technological developments will benefit the applications community. Local-area networking, shared whiteboards embedded into workstation videoconferencing tools, and the World Wide Web are all examples of network-based technological advances that have benefited the network user community. Thus, ESnet must be responsive to the needs of the programmatic applications community as well as the network research community.

Background

The ESnet Steering Committee formed an Applications Requirements Working Group to help ensure that future network requirements of the ESnet community are identified. In particular, this working group focused on applications whose demands, in the five-year time frame, are expected to exceed current network services. The working group acknowledges and appreciates the input from members of the ESnet Coordinating Committee, MICS-funded network research principal investigators, and principal investigators throughout the DOE community, including participants in the recent DOE Large Scale Networking Workshop. As we work together to realize the benefit of network advancements, we expect that the ESSC, the ESnet implementation team, and the MICS program office will supplement this document with experience acquired from DOE2000, the Network Challenged Applications program, and other contacts within the DOE research community.

The Present Role of ESnet in DOE Research

ESnet was established in 1986 to provide commonly needed network services to DOE's Energy Research programs, following a decade during which computer networking became established as an essential tool for research and each of the programs had found its own way to provide networking. A steering committee was established to provide input on networking requirements and to provide information back to the programs on developments in networking. This mission of ESnet and the composition of the ESSC were later expanded to include other significant partners within the DOE community.

The rather unexpected result of the feedback arrangement among the ESSC, the MICS program office, and the ESnet implementation team was a notable degree of coevolution of the research programs and the network. In many cases, the network has been tailored to provide for specific program requirements, and in turn the programs have been able to quickly and efficiently utilize new capabilities of the network. Network research has helped advance the network services of this community and the entire Internet. After eleven years of ESnet, there is a significant and growing degree of integration between the network and the programs. While this section of the Program Plan is primarily about future needs, it provides some brief examples of current advanced uses of the network.

Large collaborations of scientists are typical of the ESnet community. In the high energy and nuclear physics and fusion communities, for example, collaborations typically encompass hundreds of scientists who rely heavily on effective interactive, network-based communications. Collaborative authoring of research papers, computer codes, and other documents by extended groups is an early and continuing use of the network. The ability to exchange documents within minutes has transformed the effectiveness of distant collaboration by shortening the time between modification iterations by orders of magnitude. While use of email for this purpose had already provided an enormous benefit, the current embellishments of multimedia files, annotation, and real-time distributed shared workspace dramatically improve efficiency.

Subsequent improvements in bandwidth and networking tools facilitated the collaborative analysis of large data sets from multiple remote locations. This development, combined with collaborative writing, was crucial in permitting the formation of larger collaborations of scientists and enabling them to address larger scientific problems. A recent extension in this series of network enhancements, still maturing, is the ability of a scientist to participate remotely in data-taking sessions, with visual, graphical, and other data returned from the experimental site.

Remote conferencing is also important to current research work. Conferencing encompasses communications that mimic the interactions of a group of people around a conference table. Modern conferencing can provide real-time video of participants and presentations, with shared copies of documents for editing and annotation and high-quality audio communications for discussion. The requisite software must be easily installable on popular platforms, the environments created must follow consistent rules for ease of use with a reduced learning curve, and conferences must be easy to establish. Scientific collaboration, from experiment design to analysis, generally requires detailed interaction between the instrument scientists and the application scientists. This has just become possible through the use of network collaboration tools, resulting in a growing appreciation of collaborative possibilities given the right mix of tools.

These networking-based services and others have augmented the abilities and enhanced the effectiveness of the DOE scientific community. Services continue to expand and improve in response to developing needs of researchers.

Evolution of Network Requirements

While developments in the programmatic research of the ESnet community cannot be predicted with complete clarity, it is clear that the scientific community will depend increasingly on the integration of network-based services directly into the scientific environment. This ESnet Program Plan, to be available at http://www.es.net/pub/esnet-doc/esnet-program-plan/1998/index.html, provides a detailed look at future directions of these programs.

This section of the Program Plan only reports on the implications of current planning and makes a modest extrapolation for what comes later. No attempt has been made to identify explicit site connectivity or associated bandwidth requirements. Rather, the focus is on applications requiring advanced networking services. The future of network service requirements is driven not only by the forecast of future network-based applications, but also by research that will advance the performance and services of ESnet and the Internet as a whole. This section summarizes the new networking capabilities that will be needed by developments in applications and networking.

Application-driven requirements are sorted into the principal areas of:

- remote experimental operations
- distributed parallel computing
- remote/shared code development
- remote and distributed data access
- collaborative engineering
- visualization
- teleconferencing and videoconferencing.

Opportunities and requirements that are driven by anticipated service-offering advances are sorted into the network cross-cutting areas of:

- quality of service capabilities
- non-backbone connectivity
- network status and diagnostic tools for users
- scalable communications
- network and application security.

Application Areas

Remote Experimental Operations

Remote control of instruments and facilities is a very new and experimental capability. It has already found important uses in a few cases and several more trials are underway. Beginning with turbulence measurements on the Princeton Plasma Physics Laboratory's TFTR tokamak, plasma diagnostics have been run remotely from controlled fusion experiments for almost a decade. Recently this work has been extended to include proof-of-principle demonstrations of full remote control of tokamak operations, first on Alcator C-Mod at MIT and later on DIII-D at General Atomics.

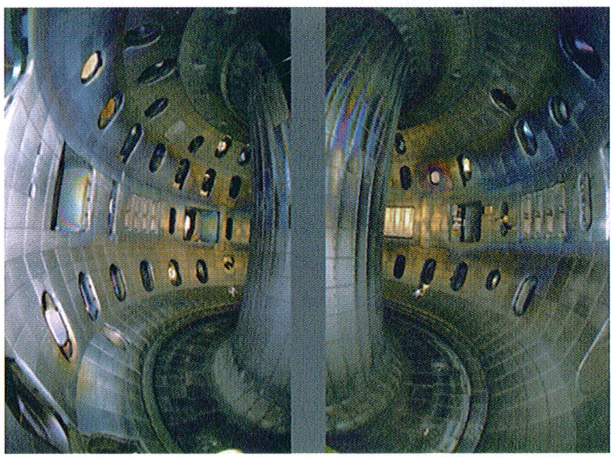

Interior of DIII-D tokamak at General Atomics in San Diego.

Developments in remote experimental operations require close cooperation between network managers and developers on the one hand and instrument builders and users on the other. Vigorous development will bring enormous benefits to research programs, including more efficient utilization of costly or unique instruments, as well as more efficient use of researchers' time. Further, the ability to correlate results from multiple instruments will greatly improve the depth of scientific understanding resulting from such measurements.

The remote control of experiments requires secure and guaranteed network transactions to ensure accurate and safeguarded control. Thus network security, authorization, and guaranteed bandwidth on demand are critical to success.

Large amounts of data are typically captured at the experiment site and are made available to the local researcher. The remote researchers must be provided with shared and distributed access to this experimental data and to shared applications. This might include access to large relational and object-oriented databases; the creation, access, and use of electronic notebooks; and visualization tools for remote data and applications performance analysis. All of these requirements will need to take place in real time, and the remote experimenter may not be located directly on an ESnet-backbone site.

Scientists involved in remote experimentation would greatly benefit from the existence of an environment (albeit virtual) similar to that at the experimental site proper. This places a demand on the network to support shared virtual environments including teleimmersion. Teleimmersion allows users to see, hear, and touch each other and a representation of their data in a simulation of the actual environment. Teleimmersion relies on virtual reality display environments at each location and high-performance network connections between locations. Resulting data flows range from streaming audio and video to virtual reality tracking data and interactive simulation updates. Additionally, diagnostics, remote access to machine status, and other related subsystems must be supported by the network with appropriate prioritization.

Authentication and security services are required at some level for all networked applications but are particularly important for experimental collaborations where expensive (and possibly hazardous) equipment may be involved.

Distributed Parallel Computing

A new computing paradigm, generating increased networking requirements as success stories spread, is distributed parallel computing. This approach to problem solving has been called by various names, including computational grids, computational nets, and simply distributed computing.

Two endpoints of distributed parallel computing are worth noting here. At one extreme, the environment would offer the sharing of a small number of distributed and unique high-performance computers. At the other extreme, the environment might consist of sharing very large numbers of underutilized desktop machines. For example, a cluster of 44 networked workstations was used to process video images containing complex three-dimensional information from the DIII-D tokamak and map the resulting data onto magnetic flux surfaces. The calculations would have taken more than a year on a single machine. Between these two extremes, a new class of distributed computing is growing, based on coupling resources on a user's desk with both local and geographically remote computing resources. As more and more sites assemble Beowulf-class computers and are willing to broker time among them, an expanded supercomputing community is forming, and associated networking requirements are expanding.

At the present time, the class of problems that lend themselves to this sort of distributed computational effort is relatively small. Algorithms that can successfully hide latency measured in tens of milliseconds instead of microseconds are still emerging from developers. However the rewards of being able to distribute problems across a meta-computer, whose resources are "free" within the brokering scheme, are so great they are being pursued by several groups. The next step in realizing such an approach is a formal, object-based computational model that allows the problem to gracefully use available resources without intrusively dominating them. Conceptually, this is a realization of the concept that the network serves as the actual backplane of the meta-computer that is composed of all the resources available on the network.

Remote/Shared Code Development

Most ESnet programs are increasingly dependent on collaborative work by researchers at distant locations. Thus, ESnet must support the collaborative development of large codes for simulation, data analysis, and other purposes. High-energy and nuclear physics collaborations typically require a centralized mechanism to control the distributed code development. As distributed computing and remote visualization become more prevalent, increased demands will be made on ESnet in this area.

Distributed code development teams can make use of desktop videoconferencing, a common code-version control system, an electronic notebook to document coding changes, common output display demonstrations, and other available communication and collaboration technologies. A common on-line code-sharing library would improve code and data access, code interconnection, and code invocation. Finally, large-scale projects, including the Spallation Neutron Source to be constructed at Oak Ridge National Laboratory, will require code development and optimization by over 100 users, well beyond the ability of present network-based tools to work effectively.

Issues such as the efficient use of distributed compute cycles, reliable asynchronous intertask communications, multicasting of data, the remote display and downloading of results, distributed task queuing, and session management will all need to be addressed if progress is to be made in this area.

Remote and Distributed Data Access

The network should enable transparent data access from remote locations for either storage or retrieval. A variety of technologies exist for this purpose—for example, distributed files systems, caching, distributed objects, and remote procedure calls—each appropriate for different applications. In general, these require a software infrastructure including an intuitive and consistent user interface, coordinated management, and other services.

Many programs using ESnet will be mounting experiments or simulation efforts that will generate enormous quantities of data. For example, the LHC (Large Hadron Collider) in Geneva, Switzerland, and RHIC (Relativistic Heavy Ion Collider) in Upton, New York, will each generate about a petabyte (million gigabytes) per year of raw data. Data rates from the Jefferson Lab accelerator in Newport News, Virginia, and the Tevatron Collider in Chicago, Illinois, are only slightly smaller. These data rates are so high that data analysis and reduction on the order of hundreds or thousands are necessary before it is feasible to move the data over the Internet. Data sets from Atmospheric Research, Basic Energy Sciences, and Fusion Energy Sciences are on the order of terabytes (thousand gigabytes). Across the ESnet community, it can be expected that in another decade at least 20 petabytes of raw data will be generated per year.

Large and rapidly growing databases of biological structure and sequence information are only tenuously connected to locally executed programs. Typically life scientists download entire databases or subsets, while others execute programs made available over the Internet which perform queries at the remote database site. This community would benefit from the existence of middleware APIs allowing programs to connect across the Internet directly to the databases, returning the appropriate files or query results to the program for immediate use.

The communities needing to use the data are distributed and will need to access the data multiple times from many different locations. In many cases, the data sets themselves will also be distributed across the network. Tools to manage and integrate views of distributed data will need to be developed. Work has begun on these issues at Berkeley Lab, Stanford Linear Accelerator Center, Brookhaven, and CERN, focused on object-oriented databases and hierarchical storage systems in support of RHIC, LHC, and the BABAR detector at SLAC.

Given such large databases, additional strategies are needed to reduce the impact of moving these data over the wide-area network. Local data reduction and data compression are two such strategies. Methods for providing users with estimates of the "costs" for accessing particular data (query estimation) may also be useful. Little has been done so far, however, to optimize network use through caching strategies. Database and network researchers will need to collaborate to solve this problem. Because generation of data is clearly increasing faster than bandwidth available on the Internet, it may be necessary to assess computational models in terms of their impact on the network.

Collaborative Engineering

Engineering has also become a collaborative effort, with design and analysis carried out by teams spread across the country and around the world. In this environment, excellent interactive communication is essential. Shared design efforts require effective tools for sharing files, displaying and annotating electronic drawings, and remote conferencing. An environment that must be virtually replicated over the network is that of a team of engineers sitting around a table piled with large, high-resolution drawings. Shared three-dimensional environments would greatly aid in visualizing complex structures and systems. Large-scale engineering codes are used in many areas of analyses. Three-dimensional thermal, structural, neutronic, and electromagnetic problems (often coupled) can often be solved only with the use of the powerful supercomputers. Engineers remote from the supercomputer centers need distributed computing tools to share the workload between local and remote systems and advanced visualization tools for analyzing the results.

Visualization

DOE scientists routinely perform computer simulations, computer modeling, and the analysis and synthesis of large amounts of experimental data, converting them into pictures or animation using sophisticated but data-intensive visualization techniques. Requirements for visualization techniques and associated data management permeate the ESnet user community, including plasma physics, climate analysis, materials, chemistry, computational fluid dynamics, combustion, DNA analysis, particle analysis, and astrophysics. Applications of such scale can be executed only in environments with large amounts of memory and processor speed. An emerging trend to address this problem is massively parallel processor (MPP)-based visualization tools, requiring connectivity with high bandwidth and low latency among the researcher's visualization environment, the MPP, and the data storage site. Given the scarcity of MPP systems, these environments are typically geographically dispersed. To achieve interactive rates, images must be delivered to the desktop at 5 to 30 frames per second, challenging the network's bandwidth and responsiveness. While interactions from the user to the MPP (typically generated by the movement of a mouse or the pressing of a switch) may not require large amounts of bandwidth, the needed response requires minimal network latency.

Similar requirements are typical of shared work spaces, immersive visualization environments including latency-intolerant haptic devices, and remote experimentation.

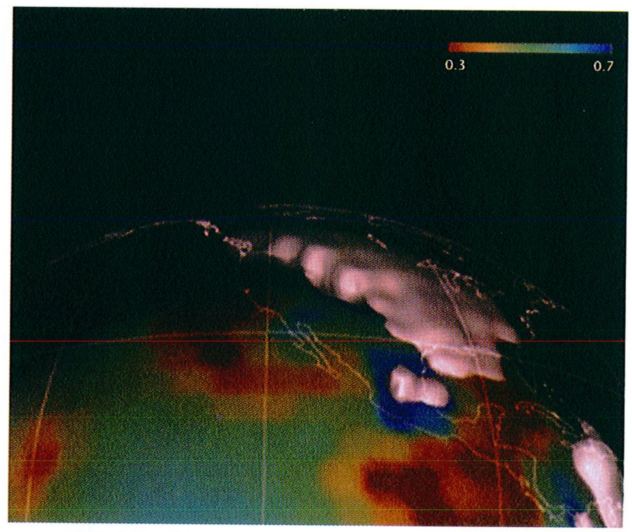

Visualization is an important tool for interpreting massive data sets. This new image of deep geological structures was revealed when innovative algorithms and NERSC's Cray T3E were used to transform seismic data from around the world into models of the three-dimensional seismic structure of the earth's crust, mantle, and core. (Donald Vasco and Osni Marques, Lawrence Berkeley National Laboratory)

Teleconferencing and Videoconferencing

The ESnet community has embraced and benefited greatly from the present availability of conference-room-based videoconferencing. Useful as it is, teleconferencing is in its infancy and needs major improvements. Advances are needed initially in ease of use, the incorporation of multimedia interactions, in the shared creation and editing of documents, and in the integration with data collection and analysis environments. The expansion of usage has increased the need for service directories, and the multiplicity of participants in any single session has created the need for floor control capabilities. Both are needed to ensure the future viability of teleconferencing. Additionally, as tools become easier to use and more integrated into the networked environment, general planning and coordination services will be needed.

Demands from the increased utilization of workstation-based videoconferencing could be enormous. The high-energy and nuclear physics communities suggest that their videoconference use in the near future might be as demanding on the network as their current data requirements—perhaps within two years as the B factories become operational in the U.S. and Japan and the Tevatron Run 2 begins at Fermilab. In the five years after that, high-energy physics will need an extensive evolution of the present conferencing system to permit effective work with gigantic data sets by extremely distributed groups of collaborators. The new capabilities required will almost certainly require evolution of relevant network protocols as well as the software at the end nodes. The ESnet community will need to have the expanded conferencing ability integrated closely with remote experimental operations, placing even more demands on the detailed operation of the network.

Conference rooms throughout the ESnet community are already fully booked with apparent unsatisfied demand. The lack of both universal interoperability and ease of use continue to pose a barrier to increased usage of this service. Commercial providers do not seem motivated to resolve this hurdle. On a positive note, standards have been recently developed in this area, and low-cost commercial implementations have created the potential for significant increased use. ESnet will need to ensure timely and supported advancements in this area.

ESnet collaborations can be expected to make heavy use of workstation-based videoconferencing. However, the existing service model is a barrier to widespread use as it only adequately supports small numbers of participants. In addition, there is a need for a complete, readily accessible directory of institutions and individuals who are accessible via this medium.

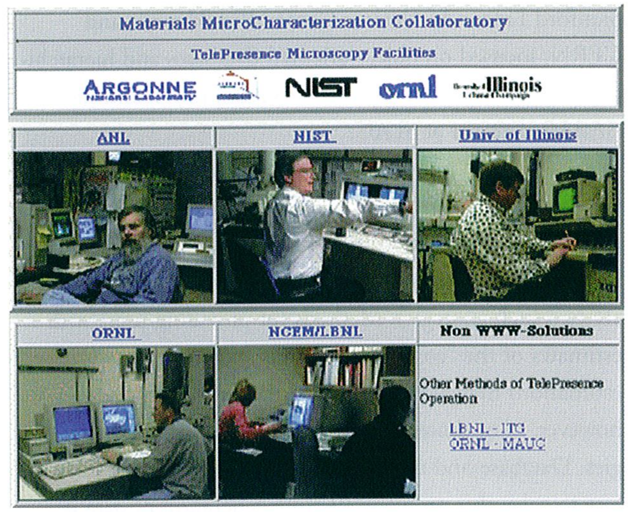

The DOE2000 Materials MicroCharacterization Collaboratory uses workstation-based videoconferencing to link scientists at five materials microcharacterization facilities in an interactive electronic laboratory.

Cross-Cutting Areas

Quality of Service Capabilities

The success of computer networking to date has encouraged the creation of increasingly demanding network-based applications and rising expectations about network performance. These new applications will require stringent limits on parameters such as latency, jitter, packet loss, and throughput.

Four factors will contribute to the growing demands on network performance:

1. Real-time applications are expected to permeate the network as collaboratory and remote experimentation benefits are realized.
2. Local-area infrastructures, whose bandwidths are currently less than ESnet, thereby serving as a bottleneck protecting ESnet, will be upgraded, allowing greater demand on ESnet services.
3. Lack of interbackbone connectivity will especially affect those users of DOE facilities who are either at university sites or not domestic.
4. Financial constraints will limit the total bandwidth available so that (1), (2), and (3) cannot be countered simply by additional bandwidth. A smarter mechanism must be found to ensure needed performance.

There currently exists no management mechanism to allocate the available resource in a manner consistent with programmatic priorities. All other DOE resources (especially user facilities such as the Advanced Light Source, the Advanced Photon Source, the Tevatron, and the National Energy Research Scientific Computing Center) have management mechanisms and implementation schemes that allow for resource allocation and protection.

Both performance and management issues require a mechanism to provide needed network resources at appropriate times. The mechanism to accomplish this is quality of service (QoS). Since different resources can be required, different QoS guarantees may be required separately or in combination. For example:

- **Network capacity or bandwidth:** An application may need a specific data rate to be successful and would need to be able to reserve that bandwidth during the session.
- **Level of packet loss:** The ability to control or guarantee a low level of lost packets may be needed by real-time or other critical applications.
- **Latency:** A low or at least constant latency may be needed for interactive sessions such as video and audio conferencing or remote control of apparatus.
- **Jitter:** Variation in packet delivery delay can be a problem for some applications (e.g., audio) and would need to be controlled within specified limits.

In some cases, QoS may be a static condition that might apply to all remote facility operation sessions. In other cases, QoS might require integration with a scheduling mechanism, so that appropriate network performance can be scheduled to coincide with other reserved resources. One can imagine that QoS would allow end-to-end prioritization of the network to coincide with scheduling time at an experiment site or the reservation of other similarly critical resources. The process of booking and then delivering a specified QoS resource is presently unsolved and is a major need of the ESnet program.

Non-Backbone Connectivity

The DOE community has benefited greatly from the architecture and authorized use policies (AUPs) of ESnet. The ESnet backbone is very responsive and generally provides more than sufficient bandwidth between a scientist and the used facility in those cases where they are both located on the backbone. However, such interconnectivity and responsiveness is not typical of the U.S. and worldwide Internet. DOE researchers who are not located on the backbone, but rather at a U.S. university or an international facility, are typically limited by the bandwidth and response time of the Internet as a whole. The communities associated with research in high-energy and nuclear physics, global climate change, the Accelerated Strategic Computing Initiative (ASCI), environmental restoration, and post-genome processing all have users and/or facilities that are not directly connected to the ESnet backbone. These require sufficient network bandwidth for dealing with very large amounts of data in experimental, computational, or analysis environments that require low network latency as well. Current Internet bandwidth and latency impair DOE science in these areas. Thus, advances are needed in end-to-end network performance and tools with an effect beyond the current borders of ESnet.

Network Status and Diagnostic Tools for Users

Information about status, configuration, and performance of the network is increasingly needed as applications rely on adequate levels of bandwidth and latency. As guaranteed performance and associated booking systems become available, booking status and associated network usage information will also be needed.

Scalable Communications

Many of the new applications described above will make use of shared data, whether the data consists of experimental data for analysis, video, audio, immersive environments, or data caching. Any one of these data transfers could require sending large quantities of data to tens of different destinations. Efficient use of network resources demands that distribution of data streams like this not be duplicated any more times than absolutely necessary. Use of a single transmission over shared portions of a route, followed by individual distribution to final destinations, is presently implemented for some applications, notably video, by multicasting. Broadcast technologies (not using the Internet) may be appropriate for simultaneous distribution of very large data sets. Continued development of these technologies will be needed, along with their integration into applications identified above.

Network and Application Security

The vision espoused in this section of the Program Plan, of new network-based paradigms such as distributed computing, virtual research groups, and telescience, depends implicitly on absolute security. This security includes both the end-to-end security of applications and, of course, security of the underlying network itself. It has come to be something of an aphorism that computers, not networks, are subject to security problems. However, the bridges, switches, routers, and domain name servers that make up the Internet are computers. They authenticate system (network) managers with username-password pairs, accept Telnet connections, and are vulnerable to exactly the same attacks that the media has reported in the computing world.

From a user's point of view, what counts is not the security of the underlying network, but the end-to-end security of applications. In this context, an application can include the remote control of an instrument or a complete research facility. Such security includes not only the protection of the bit stream, but also the security of the instrument, the facility, or even people. Providing end-to-end security at the application level involves not only securing the underlying network, it also involves authentication and certification of users or operators, and certification of application results.

Some of these issues have been solved in other contexts (e.g., training certification), and in many cases tools are beginning to appear that will allow straightforward development of solutions (such as public key encryption). But in general, complete solutions do not exist now and cannot be purchased as turnkey packages from commercial vendors. The commercial marketplace is investing in the development of secure electronic commerce (so people can send credit card numbers over the Web), but not to guarantee the results of a physics calculation. Thus, there exists a gap (recognized by DOE2000) between simple encryption of credit card numbers at one extreme, and the protection of multi-user, multi-system, real-time sessions at the other. This gap includes requirements for certification and authentication of users across administrative domains, protection and certification of shared information or results, and very high-speed encryption of real-time data streams at OC12 speeds or higher.

Much of the work needed is summarized in the paper "Research & Development Priorities for Communications and Information Infrastructure Assurance" by Huntman, Jacobsen, Johnson, Mansur, and Baily (available at http://www-itg.lbl.gov/security/Publications/C+I_Report.html). The paper provides rough estimates of the levels of research needed in 13 topical areas: (1) characterization and notification of threats; (2) detection, analysis, and prevention; (3) definition of security architectures; (4) response, recovery, and reconstitution; (5) advanced concepts and theory; (6) management of information protection; (7) characterization of infrastructure required for minimum essential services; (8) valuation of information; (9) indication and warning; (10) cost-benefit analysis; (11) modeling and simulation; (12) risk management; and (13) encryption technologies.

Network Research and Development

The MICS-funded network research and development efforts are intended to create and/or enable new high performance networking applications. Examples of their efforts to ensure scalable networking include the improvement of network protocols and router algorithms; the creation of tools that model, measure, and analyze network traffic; technology that can guarantee bandwidth (e.g., quality of service) and associated management tools; multicasting of data; security; and innovative techniques for the efficient handling of World Wide Web-related traffic. ESnet is motivated to support these efforts because such R&D underpins the future success of networking, and the R&D program itself is a member of the user community that ESnet is mandated to serve.

This situation presents a unique set of challenges to the ESnet provider. On the one hand, users of the production aspects of ESnet do not wish to have their network infrastructure disturbed by such R&D activities. On the other hand, network researchers must have access (at some point) to a large network with real users as the ultimate testbed for their efforts. Such an environment is critical to the future of ESnet and the entire community that it serves. A full discussion of issues related to supporting production and research traffic (and one possible solution) can be found in "MORPHnet" by Aiken, Carlson, Foster, Kuhfuss, Stevens, and Winkler at http://www.anl.gov/ECT/Public/research/morphnet.html. Such a testbed is an important requirement.

In the end, research and development must result in products that work in the scientific environment of the ESnet user community. Such products must be reasonably easy to use, perhaps appearing almost shrink-wrapped. Such tools must be integrated into the full scientific research environment and be easily maintainable.

A Collaborative Approach

This section has discussed the need for a variety of complex and advanced networking services. In order to ensure their timely introduction and ease of use by the scientific community, a collaborative approach to development and implementation is needed. Most of the solutions will require close collaboration among the ESnet staff, the scientists using the network, and the network research and development staff. Of the applications and infrastructure enhancements discussed above, this collaborative approach to development will be particularly needed for remote experimental operations, remote/shared code development, visualization, teleconferencing/videoconferencing, quality of service, and network status and diagnostic tools for users.

Without a collaborative and tightly coupled approach, application and network infrastructure developments will not be coordinated, resulting in less than synergistic efforts and lost time. ESnet, as a network run for DOE research programs in close consultation with the user programs, is ideally positioned to ensure the needed collaborative developments and to work iteratively with the programmatic users to ensure that the users get the maximum value from the network and the improved functionality they need as soon as possible.

Close cooperation between network researchers, network service providers, and the ESnet user community will continue to be central to the success of the DOE networking program and the community that it benefits.

Acknowledgment

This section is based on a white paper on advanced applications written for the ESSC by Martin Greenwald, Sandy Merola, Larry Price, and Bill Wing.

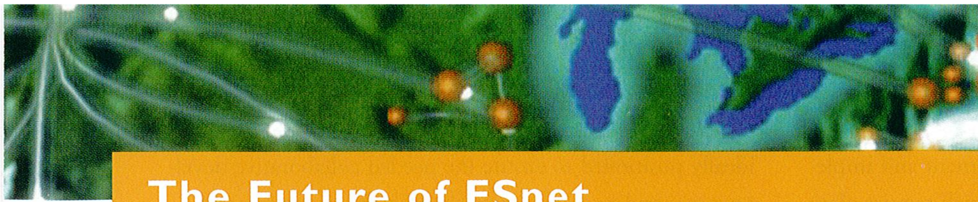

The Future of ESnet

ESnet faces some significant challenges in the future:

- Of paramount importance is the ability to continue providing, in the face of limited financial resources, an increasing level of service to users, generated by new applications.
- The relationship and interactions with emerging initiatives need to be fostered such that the current user base is not adversely affected.
- The need for expensive and scarce international connectivity will increase, because of the reliance on foreign research centers, e.g., the Large Hadron Collider at CERN.
- ESnet needs to foster emerging network technologies and continue its active network research program while maintaining a stable service-oriented network.

The past success of ESnet in meeting challenges is assurance that ESnet will meet the challenges of the future. The past success is due, in large measure, to the interaction of five key elements of ESnet—Federal program management, ESnet staff, network researchers, applications developers, and users—which is unusually effective and has led to a cooperative problem-solving atmosphere. The successful interaction of these elements make ESnet the most respected network supporting scientific research in the world. Thus, in spite of the difficulty of predicting tomorrow's technology, it can be stated with confidence that ESnet will keep abreast and provide leading-edge infrastructure.

In a rapidly changing techno-political environment, there are many forces that challenge and question the continuation of standard operation procedures of the past. The challenge for ESnet is to continue those practices and policies of the past that have distinguished it, so that support from the users and the funding sources will continue. The synergistic relationship among ESnet users, operations, and management is difficult to quantify, but is evident in the degree to which the mission of ESnet is fulfilled.

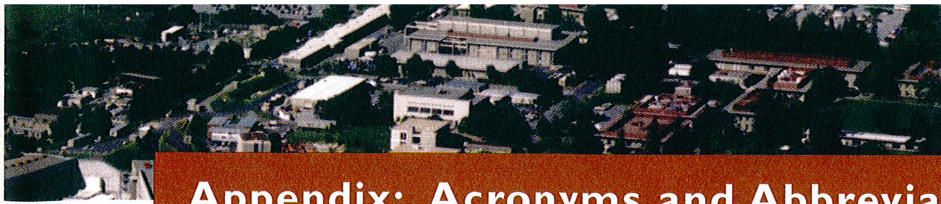

Appendix: Acronyms and Abbreviations

ACL	Advanced Computing Laboratory, LANL
ACTS	Advanced Computational Testing and Simulation
AFS	Andrew File System
AGEX	aboveground experiments
AGS	Alternating Gradient Synchrotron, BNL
ANL	Argonne National Laboratory
ARM	Atmospheric Radiation Measurement Program
ARPA	Advanced Research Projects Agency
ASCI	Accelerated Strategic Computing Initiative
ATLAS	Argonne Tandem Linac Accelerator
ATM	Asynchronous Transfer Mode
AUP	authorized use policy
BABAR	B/B-bar detector, SLAC
BEPC	Beijing Electron-Positron Collider
BER	Biological and Environmental Research
BES	Basic Energy Sciences
BITnet	"Because It's Time" Network
BNL	Brookhaven National Laboratory
bps	bits per second
CADD	computer-aided design and drafting
CART	Cloud and Radiation Testbed
CBR	constant bit rate
CCS	Center for Computational Sciences, ORNL
CDIAC	Carbon Dioxide Information Analysis Center
CEBAF	Continuous Electron Beam Accelerator Facility (now JLab)
CERN	European Organization for Nuclear Research
CESR	Cornell Electron Storage Ring
CHAMMP	Computer Hardware, Advanced Mathematics, and Model Physics
CRADA	Cooperative Research and Development Agreement
CTBT	Comprehensive Test Ban Treaty
DANTE	Delivery of Advanced Network Technology to Europe Ltd.
DARPA	Defense Advanced Research Projects Agency
DCC	Diesel Combustion Collaboratory
DESY	German Electron Synchrotron Laboratory
DOD-IP	Department of Defense Internet Protocol
DOE	Department of Energy
DP	Office of Defense Programs
DPSS	Distributed Parallel Storage System
DSIT	Data and Science Integration Team
ECCE'	Extensible Computational Chemistry Environment
EE	Office of Energy Efficiency and Renewable Energy
EIA	Energy Information Agency
EMSL	William R. Wiley Environmental Molecular Sciences Laboratory
EPR	electronic paramagnetic resonance
ER	Energy Research
ERSUG	Energy Research Supercomputer Users Group
ESCC	ESnet Site Coordinating Committee
ESnet	Energy Sciences Network
ESSC	ESnet Steering Committee
ETP	Environmental Technology Partnerships
EXERSUG	Executive Committee, Energy Research Supercomputer Users Group
FDDI	fiber distributed data interface
FES	Fusion Energy Sciences
FIX-E	Federal Interagency eXchange-East
FIX-W	Federal Interagency eXchange-West
FNAL	Fermi National Accelerator Laboratory
FTICR	Fourier transform ion cyclotron resonance mass spectrometer
GA	General Atomics
Gbps	giga (billion) bits per second
GC	Grand Challenge
GCM	general circulation model
HEP	High Energy Physics
HEPAP	High Energy Physics Advisory Panel

HEPnet	High Energy Physics Network
HERA	Hadron Elektron Ring Anlage, DESY
HIPPI	high performance parallel interface
HPC	high performance computing
HPCC	High Performance Computing and Communications program
HPCRC	High Performance Computing Research Center
HR	Human Resources and Administration
HRC	HEPnet Review Committee
ICFA-NTF	International Committee on Future Accelerators-Network Task Force
IEEE/NPSS	Institute of Electrical and Electronics Engineers, Nuclear and Plasma Sciences Society
IETF	Internet Engineering Task Force
IHEP	Institute for High Energy Physics, Protvino, Russia
IHEP	Institute of High Energy Physics, Beijing, China
INSP	International Nuclear Safety Program
I/O	input/output
IP	Internet protocol
ISMA	Internet Statistics and Measurements Analysis
ITER	International Thermonuclear Experimental Reactor
JET	Joint Engineering Team
JINR	Joint Institute for Nuclear Research
JLab	Thomas Jefferson National Accelerator Facility (formerly CEBAF)
kbps	kilo (thousand) bits per second
KEK	Japanese National Laboratory for High Energy Physics
LAN	local-area network
LANL	Los Alamos National Laboratory
LBNL	Ernest Orlando Lawrence Berkeley National Laboratory
LEP	Large Electron-Positron Collider, CERN
LHC	Large Hadron Collider, CERN
LLNL	Lawrence Livermore National Laboratory
Mbone	Multicast backbone
Mbps	mega (million) bits per second
MFEnet	Magnetic Fusion Energy Network
MICS	Mathematical, Information, and Computational Sciences Division
MIPS	million instructions per second
MIT	Massachusetts Institute of Technology
MMC	Materials MicroCharacterization Collaboratory
MOU	Memorandum of Understanding
MPP	massively parallel processor
MSCF	Molecular Science Computing Facility, EMSL
NABIR	Natural and Accelerated Bioremediation Research
NANOG	North American Network Operators Group
NAS	National Academy of Sciences
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NE	Office of Nuclear Energy
NERSC	National Energy Research Scientific Computing Center, LBNL (formerly National Energy Research Supercomputer Center and NMFECC at LLNL)
NFS	Network File System
NGI	Next Generation Internet
NIH	National Institutes of Health
NII	National Information Infrastructure
NIST	National Institute of Standards and Technology
NMFC	National Magnetic Fusion Collaboratory
NMFECC	National Magnetic Fusion Energy Computer Center (now NERSC)
NMR	nuclear magnetic resonance
NREN	National Research and Education Network
NSF	National Science Foundation
NTTP	Numerical Tokamak Turbulence Project
OBER	Office of Biological and Environmental Research
OBES	Office of Basic Energy Sciences
OCTR	Office of Computational and Technology Research
OFES	Office of Fusion Energy Sciences
ORNL	Oak Ridge National Laboratory
PEP-II	Positron-Electron Project II, SLAC
PNNL	Pacific Northwest National Laboratory
PPPL	Princeton Plasma Physics Laboratory

QoS	quality of service
REE	Remote Experimental Environment
RHIC	Relativistic Heavy Ion Collider, BNL
RISC	reduced instruction set computer
RTG	radioisotope thermoelectric generator
SBSS	Science-Based Stockpile Stewardship
SDR	Soviet-designed reactor
SLAC	Stanford Linear Accelerator Center
SLC	Stanford Linear Collider
SNL	Sandia National Laboratories
SOC	self-organized criticality
SSC	Superconducting Super Collider
TCP	transmission control protocol
TFTR	Tokamak Fusion Test Reactor
vBNS	very high speed Backbone Network Service
VCS	ESnet Video Conferencing Service
WAN	wide-area network
WDC	World Data Centers

Available to DOE and DOE Contractors from
the Office of Scientific and Technical Information:
DOE Technical Information Services
P.O. Box 62,
Oak Ridge, TN 37830
(423) 576-8401
http://www.doe.gov/html/osti/contact.html

Available to the public from
the National Technical Information Service:
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161
(703) 487-4650
http://www.ntis.gov/ordering.htm

Available on the World Wide Web at:
http://www.es.net/pub/esnet-doc/esnet-program-plan/1998/index.html

Visit ESnet on the Web at:
http://www.es.net/

Printed on recycled paper

U. S. Government Printing Office: 1998 433-460/60069

U.S. Department of Energy
Office of Energy Research
Mathematical, Information,
and Computational Sciences Division
Germantown, MD 20874