DOE/SC-0031
U.S. Department of Energy

ESnet

Program Plan 2001

ESnet

The cover shows a simulation of the electric potentials in a magnetically confined plasma under conditions similar to what would occur in a nuclear fusion reactor. If the central temperature becomes too high, the plasma becomes turbulent. In that state, it is harder for the strong magnetic field around the plasma to provide the thermal insulation required for the fusion reaction to occur. The image shows the potentials just before the transition to turbulence. Such simulations help researchers discover the conditions that will make the plasma confinement most effective.

Image courtesy of Scott Parker, University of Colorado

DOE/SC-0031
March 2001

ESnet

Program Plan 2001

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

Energy Sciences Network Steering Committee (ESSC) 2001

Basic Energy Sciences

Jeffrey A. Nichols
Technical Group Leader
High Performance
Computational Chemistry
Pacific Northwest National Laboratory
P.O. Box 999, MS K1-96
Richland, WA 99352
E-mail: Jeff.Nichols@pnl.gov

Nestor J. Zaluzec
Material Sciences Division
Argonne National Laboratory
9700 South Cass Avenue
Argonne, IL 60439
E-mail: zaluzec@aaem.amc.anl.gov

Defense Programs

Pete Dean
Sandia National Laboratory
P.O. Box 969
Mail Stop 9011
Livermore, CA 94550-9011
E-mail: pwdean@sandia.gov

Fusion Energy

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

Stephen C. Jardin
Principal Research Physicist
Princeton Plasma Physics Laboratory
Department of Physics
P.O. Box 451
Princeton, NJ 08543-0451
E-mail: jardin@pppl.gov

Health and Environmental Research

Jimmy W. Voyles
Manager, ARM Data and
Science Integration
Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA 99352
E-mail: jimmy.voyles@arm.gov

Reinhold C. Mann
Director, Life Sciences Division
Oak Ridge National Laboratory
1000 Bethel Valley Road
P.O. Box 2008
Oak Ridge, TN 37831-6124
E-mail: mannrc@ornl.gov

High Energy Physics

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

International Nuclear Safety

Greg Chartrand
Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA 99352
E-mail: greg@pnl.gov

Mathematical, Information, and Computational Sciences

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

Charlie Catlett
Senior Fellow
Mathematics and Computer
Science Division
Argonne National Laboratory
3315 Digital Computer Laboratory
1304 W. Springfield
Urbana, IL 61801
E-mail: catlett@mcs.anl.gov

Nuclear Physics

Chip Watson
Thomas Jefferson National
Accelerator Facility
MS 16A
12000 Jefferson Ave.
Newport News, VA 23606
E-mail: Chip.Watson@jlab.org

ESnet Coordinating Committee (ESCC)

William R. Wing
Oak Ridge National Laboratory
P.O. Box 2008, MS6144
Oak Ridge, TN 37831-6144
E-mail: wrw@ornl.gov

Energy Sciences Network (ESnet)

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

February 5, 2001

Dr. James F. Decker, Acting Director, DOE Office of Science
Dr. C. Edward Oliver, Associate Director for Advanced Scientific Computing Research
Mr. George Seweryniak, ESnet Program Manager

The Energy Sciences Network (ESnet) Steering Committee is pleased to submit to you the 2001 ESnet Program Plan. If you compare it with the previous Program Plan, published in 1998, you will see that many changes have taken place. The programs of the Office of Science and other parts of DOE that rely on ESnet have advanced in many ways. This is especially so with respect to the quantities of data produced and the need for geographically separated researchers to share and work collaboratively on those data. The result is that ESnet is more than ever at the heart of the DOE research enterprise. ESnet has responded by expanding both high-speed connectivity and the advanced networking services needed for collaboratories; emerging grids for computing, storage, and visualization; and related tools used to enhance the productivity of DOE research.

The ability of ESnet to respond quickly and effectively to the evolving needs of DOE missions is strengthened by the strong cooperation between ESnet providers and users. The Steering Committee has been a part of the ESnet process from the beginning. It is composed primarily of representatives of the programs that use ESnet, in addition to providers of ESnet services. It is charged with identifying and prioritizing network requirements and with reviewing implementation plans and resource allocation. Standing subcommittees and flexible task forces and working groups ensure that members of the user community are strongly involved with the evolution of ESnet.

This Program Plan was requested by the Mathematical, Information, and Computational Sciences Division of the DOE Office of Science. It describes the current configuration of ESnet and the major directions that are foreseen for its development in the next few years. A large fraction of the Plan is devoted to the programs served by ESnet and the requirements they foresee, based on the future directions planned for their programs. Coordinated by the Steering Committee, the document is itself the result of a broad collaborative process with contributions from many principal investigators across the Office of Science and the Department of Energy.

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

Table of Contents

CORE PROGRAM REQUIREMENTS - Chapter 3

Executive Summary

ESnet and its predecessors in the U.S. Department of Energy (DOE) have provided leading-edge networking for DOE missions since the emergence of wide-area networking. The evolution of scientific modeling and experimental science, combined with state-of-the-art computing technologies, has increased the reliance of the scientific community on advanced networking bandwidth and services. In fact, without the capabilities of ESnet, the scientific mission of the DOE could not be achieved.

Over the years, we have seen vast advances in available network bandwidth and services, as well as the emergence of the global Internet, which is modeled on ESnet and other Internet pioneers. The constant amid all this change has been the presence of mission-driven requirements for advanced networking capabilities not available from the commodity market. The most demanding requirements have come from DOE's science programs in what is now the Office of Science, although several other parts of the Department have found ESnet to be important in their work.

New capabilities and services developed for ESnet support new scientific opportunities and help DOE science provide the best return on taxpayer investment. ESnet has proved to be truly an enabling technology and service that has enhanced the capabilities and successes of DOE's science programs across the board. The following are some examples, which are described in more detail in this program plan.

- *Instant communications and the ability to confer quickly about large data sets have enabled High Energy Physics to mount broad collaborations that unify the best experts around the world, for example, to find the top quark, discover neutrino mass, and search for the Higgs boson.*
- *The extremely capable network and the collaboratory services of ESnet have enabled Basic Energy Sciences to provide broad access to specialized instruments, such as microscopes and synchrotron light sources. The availability of these instruments has increased the speed and significance of discoveries in this program.*
- *By providing a means to link the control rooms of major fusion experiments to remote researchers via streaming audiovisual and high-performance data links, ESnet enables these experiments to operate as shared national facilities.*
- *Progress in networking bandwidth, tools, and middleware has enabled (and is motivated by) computing and computational sciences to advance scientific computing tools at an unprecedented rate.*

In this program plan, we lay out the requirements of the programs that rely on ESnet and describe our plans for meeting these requirements during the next three years. It is important to note that in the fast-changing world of science and technology, no program can describe with precision what it will need even three years from now. So an overarching requirement from all the programs is that ESnet must be flexible and adaptable to meet emerging needs.

ESnet will use two strategies to meet the challenge of evolving requirements. First, we will continue our uniquely successful tradition of user-level committees that provide input on all aspects of ESnet. There are now three such committees:

- *ESnet Steering Committee (ESSC), which defines requirements for ESnet and gives it general guidance.*
- *ESnet Coordinating Committee (ESCC), which provides technical expertise and liaison with the ESnet sites.*
- *ESnet Research Support Committee (ESRSC), a new committee established in 2000, which coordinates the use of advanced ESnet resources to support high-performance network research and testing.*

Second, results of new computing research will be rapidly incorporated into the operational network. This approach is particularly important because the network services envisaged in scientific research plans are so advanced that the technological path for providing them has yet to be defined. Most of the fields represented in this program plan indicate a need for "grids." These range from now-emerging computational grids needed by computationally intensive fields (e.g., molecular orbital calculations and many modeling and simulation tasks) to the proposed virtual data grids needed by data-intensive fields (e.g., collider detectors and atmospheric measurements).

Meeting these requirements will take a challenging interplay among the operation of the production network, computing and network research, and testing and deployment of new capabilities on a rapid schedule. Close contact with the research and technical communities is essential. Chapter 2 of this program plan describes several other emerging technologies that will be important to supporting future DOE research.

Continuation of the indispensable benefits of ESnet to the DOE programs will depend on efforts in the following areas and will require continued increases in funding of roughly $1.5 million to $2 million a year.

- *ESnet will need to expand performance levels to meet the expanding program requirements. Network rate capability is expected to need exponential growth of a factor of two per year, on the basis of past experience and programmatic predictions.*
- *ESnet must also expand and enhance the services that allow programmatic users to use the network for collaboration.*
- *Finally, because most programs are relying on grid-type distribution of computing and data resources for the next advance in their research, ESnet will need to work closely with the research and development projects aimed at creating these capabilities. In this regard, the new focus on ESnet testbeds and other uses of the network to support research in network capabilities will be vital to the future success of ESnet and DOE.*

1 ESnet Plan

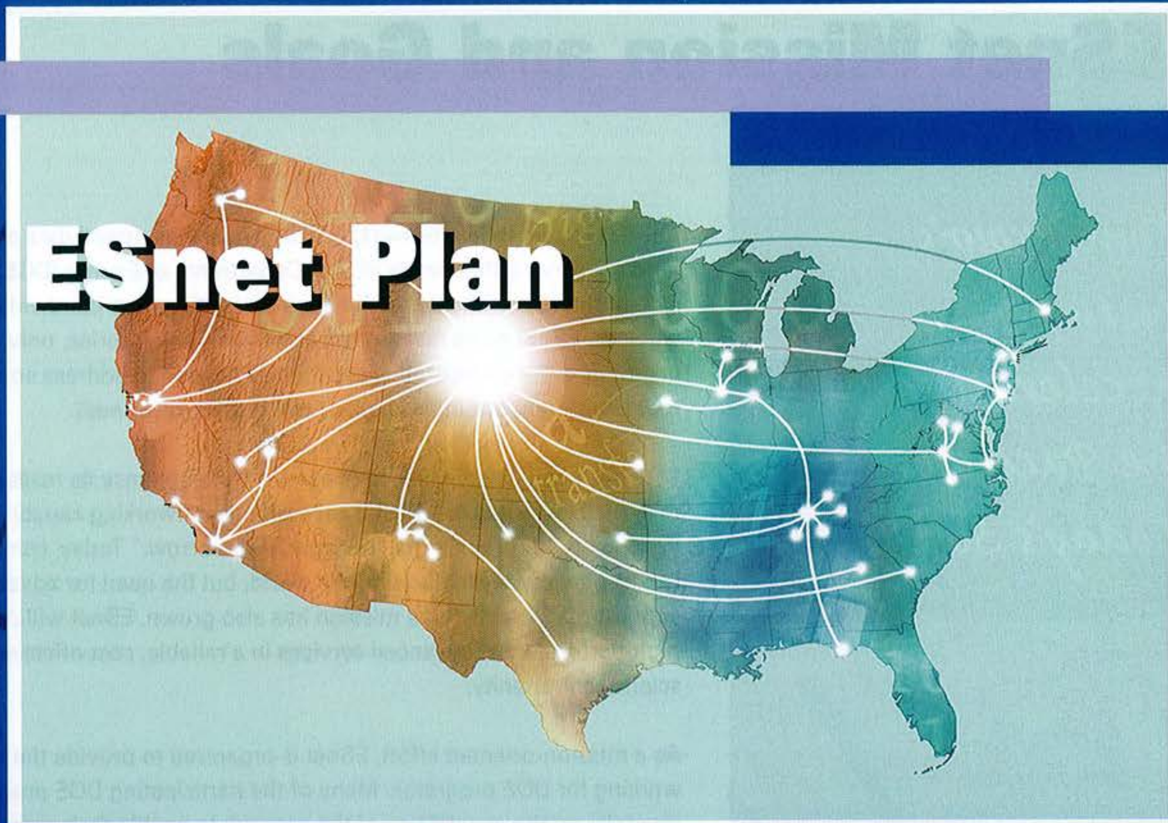

ESnet Mission and Goals

Tomorrow's Network Accelerating the Pace of Today's Science

The Energy Sciences Network, or ESnet, is a high-speed data communications network serving thousands of U.S. Department of Energy (DOE) scientists and their collaborators worldwide. A pioneer in providing high-bandwidth, reliable network services, ESnet gives researchers at national laboratories, universities, and other institutions the collaborative capabilities needed to address some of the world's most important scientific challenges *(http://www.es.net/)*.

DOE has been using advanced networking to enhance its research activities since the mid-1970s (see pages 13-15). Leading-edge networking capabilities are adopted—or developed—to provide the "network of tomorrow." Today, commodity networking is part of everyday life around the world, but the need for advanced networking to support the Department's mission has also grown. ESnet will continue to provide high-performance, advanced services in a reliable, cost-effective manner to the DOE science community.

As a mission-oriented effort, 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 instruments—all of which can be summarized as research enhancement.

This vital connection has led ESnet to pursue the following goals, the implementation of which is described in this program plan:

- Reliable, production-quality network services with capabilities based on leading-edge technology.
- Close coupling to the Department's programmatic requirements.
- Ongoing improvements in network services and related applications targeted at the rapidly evolving and growing needs of the programs.
- Highly leveraged interaction and coordination with ESnet sites to optimize service, performance, and resources.
- Effective interagency and international coordination and cooperation.

ESnet is vital to DOE research

As an agency that is at the core of fostering and funding research in the United States, DOE finds it imperative to implement technology and services that will

- *Accelerate the pace of scientific research.*
- *Support the global agency research community.*
- *Allow for effective and remote use of DOE's unique facilities.*
- *Maximize the effectiveness of funding for agency-wide benefit.*

Computer networking is a proven technology that can help accomplish these goals, and ESnet has a leading role worldwide in supporting the scientific research community. ESnet is therefore a vital ingredient for successfully achieving the Department's mission in the next decade.

ESnet Today

The Network at a Glance

ESnet provides advanced computer networking for the DOE science mission and other DOE missions. Managed and operated by the ESnet staff at Lawrence Berkeley National Laboratory, ESnet currently provides direct connections to more than 30 major DOE sites at speeds up to 622 megabits per second (Mbps). Connectivity to the global Internet is maintained through interconnection ("peering") arrangements with more than 100 other Internet service providers.

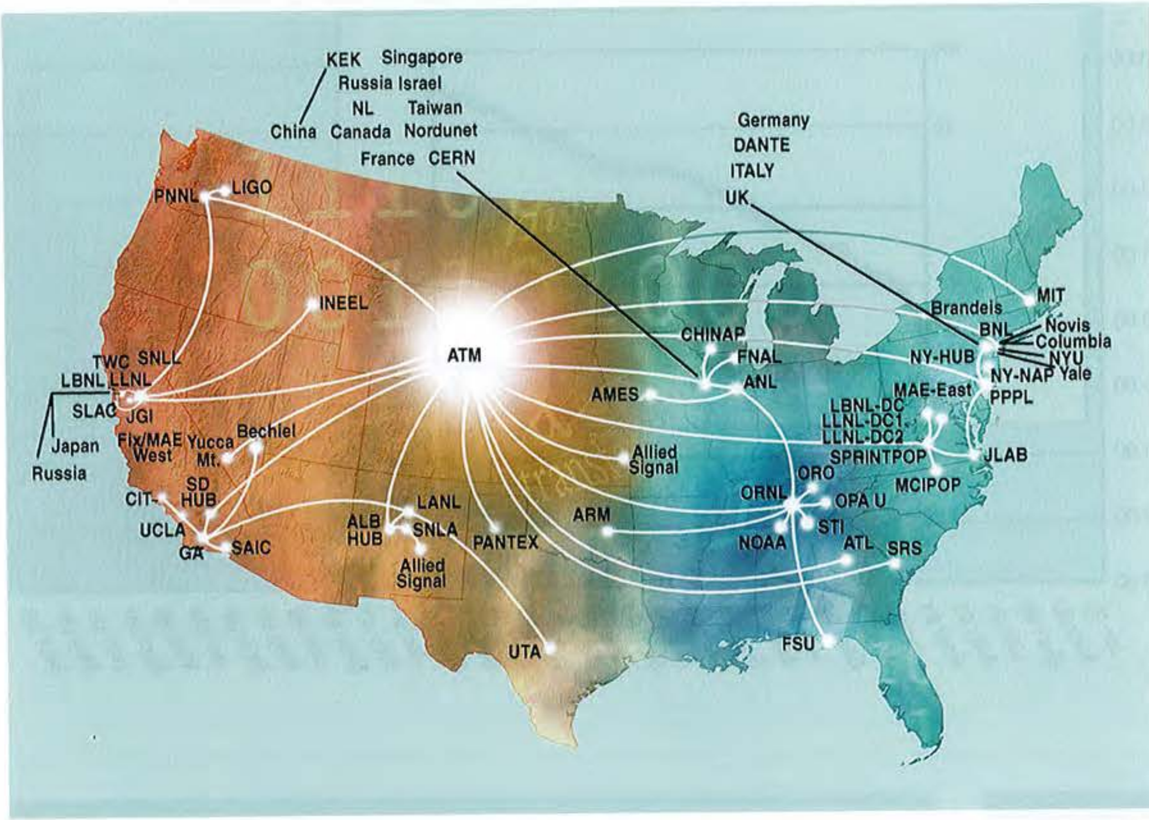

ESnet connectivity as of summer 2000, showing asynchronous transfer mode (ATM) links

Funded principally by DOE's Office of Science, ESnet allows DOE scientists to use unique DOE research facilities and computing resources independent of time and location with state-of-the-art performance levels. ESnet supports science in ways that extend from the mundane, such as communication via e-mail, to the data- and bandwidth-intensive, such as rapid distribution of enormous volumes of physics data to researchers around the world. Beyond data transfer, ESnet provides capabilities that are truly amazing, such as "telepresence" and remote operation of equipment, in which physically separate groups interact with each other and with unique instruments as if they were in the same room.

Performance Growth

Since its inception in the mid-1980s, ESnet performance levels have grown with user demands by a factor of more than 10,000. The amount of traffic accepted by the network from researchers has grown by 100% every year since 1990 (see figure). Under current projections, performance may grow by another factor of more than 1,000 within the next five years. This growth will be necessary to support DOE's role in proposed interagency programs. The new performance levels may include the ability to network at terabit-per-second (Tbps) levels (1 million Mbps). ESnet is already planning an architecture and approach to accommodate such traffic.

Monthly Internet protocol (IP) traffic over ESnet, in terabytes. Inset graph shows the same data on a logarithmic scale, which shows the steady exponential increase.

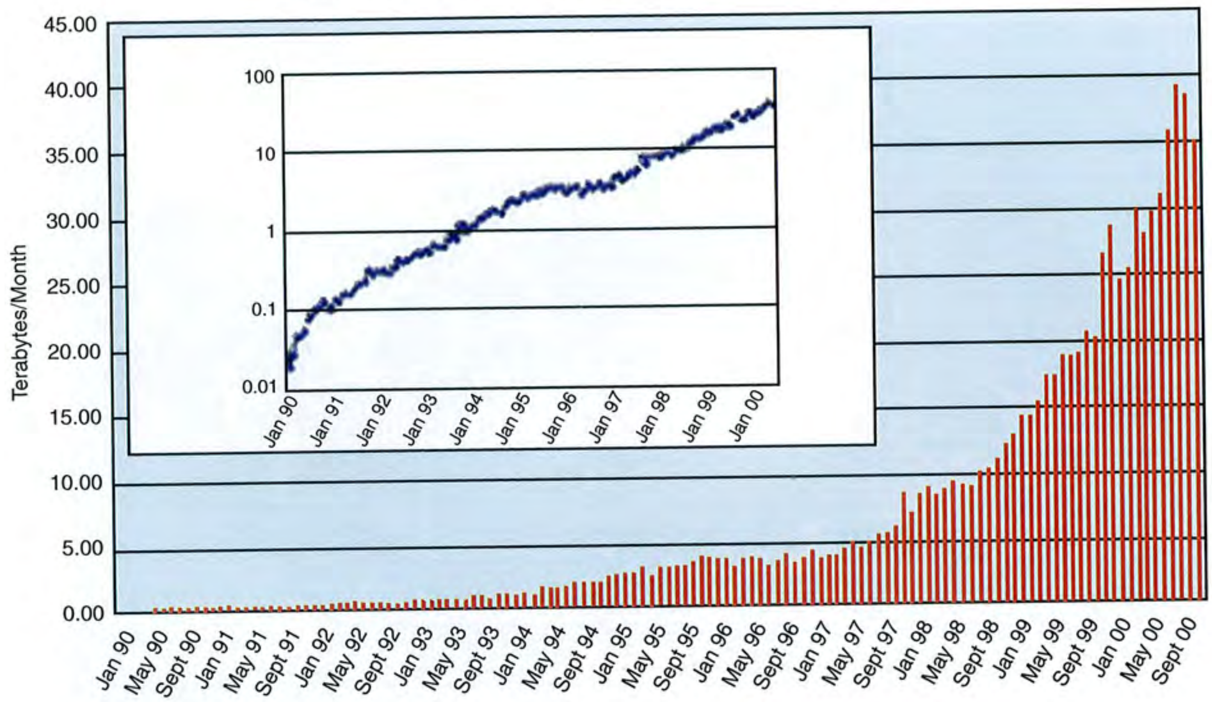

Community Involvement

ESnet derives its effectiveness from the extensive cooperation it enjoys with its user community. It is one of the most widely based and successful cooperative efforts within DOE, perhaps within the federal government. The extensive participation by technical and scientific personnel on a DOE-wide basis results in extremely effective leveraging of both effort and resources. A recent formal review of the ESnet program gave it a rating of "outstanding" and commented on its high degree of cost-effectiveness.

The ESnet Project enjoys an excellent working relationship with both its technical and program user communities. This relationship is maintained through three user committees:

- The ESnet Steering Committee (ESSC) deals with ESnet's strategy, policy, operational requirements, and priorities as established by principal investigators representing DOE program areas.

- The ESnet Coordinating Committee (ESCC) deals with the associated technical issues affecting the sites and the backbone.
- The ESnet Research Support Committee (ESRSC) deals with the requirements and technical issues related to ESnet testbed activities.

These three committees sponsor special-interest task forces and working groups that study issues, establish consensus, and share recommendations and information.

An essential element of DOE research is collaboration among teams of researchers located around the world. ESnet enhances the effectiveness of these scientists' work by providing a rich interconnectivity to the "outside" world. The network includes interconnections to many other U.S. networks, as well as several direct connections to international sites and networks. A recent emphasis has been enhancing interconnectivity to U.S. universities with "peering" interconnects to Abilene (the Internet2 backbone network) at speeds up to OC-12 (622 Mbps).

Research

The user-driven, collaborative framework of ESnet, with its ongoing combination of reliable services coupled with state-of-the-art capabilities, uniquely positions it to contribute to the development of leading-edge technologies. ESnet is also strategically positioned to participate in interagency, national, and international research and development (R&D) projects. Specific areas of current research and advanced technology include Internet protocol (IP) version 6 (IPv6), asynchronous transfer mode (ATM), streaming video, multicasting, virtual private networks, dense wave division multiplexing (DWDM), and IP differentiated services.

Large data files shared with international collaborators through ESnet

The BaBar detector, built by a large international team of physicists, detects the B/B-bar system of mesons produced by the PEP-II collider at the Stanford Linear Accelerator Center (SLAC). It consists of a silicon vertex detector, a drift chamber, a particle identification system, a cesium iodide electromagnetic calorimeter, and a magnet with an instrumented flux return.

This complex device produces an enormous amount of data requiring considerable computing resources for analysis. The data sets are being sent across the network to be analyzed at several collaborating institutions in Europe.

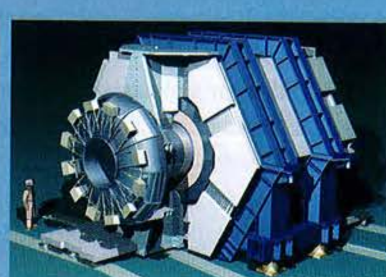

Remote collaboration and experimentation speeds development of thin-film technologies

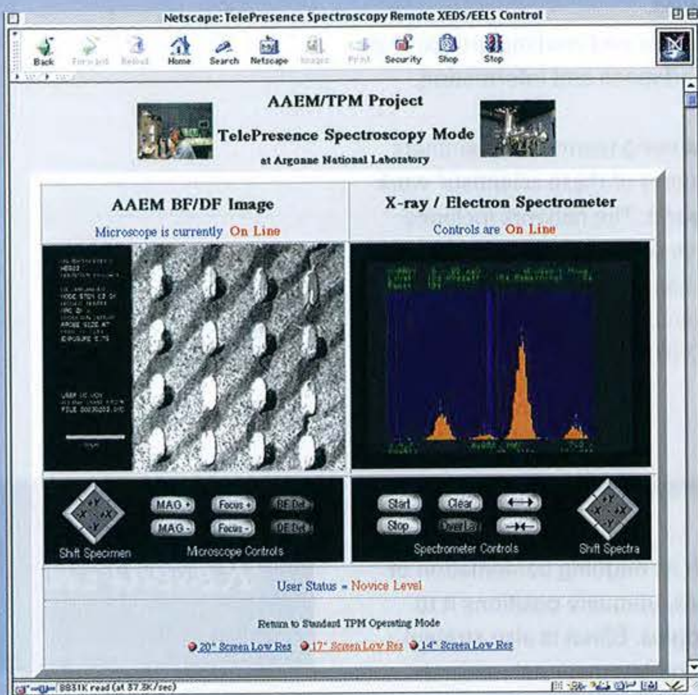

This image is from a study of the way magnetic domains in a thin film are influenced by its artificially tailored microstructures (for details, see http://tpm.amc.anl.gov/mmc/Pubs/NanoMagnet/NanoMagnetic1.html).

To accelerate the process by which science and innovation are accomplished, "telepresence" collaboration tools are used to bring together different kinds of scientists from geographically separated research groups. Such tools are being used for collaborative imaging and spectroscopy of thin films (see figure). Scientists can now control the growth of thin films and characterize these materials at the atomic level, and as a result they can produce tailor-made complex nanostructures with properties not found in nature. Researchers are using these novel materials to answer fundamental scientific questions and to provide society with new technologies that offer a better quality of life.

Other Services

The ESnet Project provides additional services, beyond those mentioned above, to support the research efforts in DOE. The flagship effort in this area is the ESnet digital collaboration services, or DCS. The DCS is a collection of collaboration services and tools brought together in an integrated manner and designed for ease of use. The tools include a multi-point control unit for videoconferencing among several participants, an audio bridge, and a data conferencing tool. All components can be reserved via a Web-based reservation tool, and they can be combined in ways that enhance the effectiveness of each one. For example, a user can reserve both the control unit and the audio bridge for a videoconference. With this arrangement, a researcher who is traveling and does not have access to video equipment, for example, can still participate in the conference by voice.

3 Core Program Requirements

Mathematical, Information, and Computational Sciences

ESnet

Program Description

Facilities

Science and technology research within DOE involves a variety of analysis, modeling, simulation, and prediction activities based on both experimental and simulation data sets. The Mathematics, Information, and Computational Sciences (MICS) program supports these activities through advanced fundamental research in applied mathematics and computer science, concentrating on areas that are essential to the work of DOE researchers and technologists. MICS also operates high-performance computing, networking, and related capabilities for the DOE research community in general. These facilities collectively provide opportunities both to develop technology at scale and to make new technologies available to researchers in a timely fashion. ESnet, operated by MICS, is central to this technology creation and early deployment pipeline.

The MICS program also operates advanced computing research facilities that support computational science and engineering and provide an environment for creating and proving new computing, information, and communications technologies. These facilities are located at

- Argonne National Laboratory (Mathematics and Computer Science Division).
- Lawrence Berkeley National Laboratory (National Energy Research Scientific Computing Center).
- Los Alamos National Laboratory (Advanced Computer Laboratory).
- Oak Ridge National Laboratory.

Research and Development Focus

The computational tools and technologies that are available to energy researchers continue to improve at tremendous rates. During the 1980s, computing power available to researchers increased by three orders of magnitude. The 1990s brought even greater improvements. In 1992, a "supercomputer" in DOE was measured in gigaflops: a $35-million, 1024-processor CM-5 was capable of 130 gigaflops. By contrast, in 2001 it is possible to build a system with the same 130 gigaflops and advanced graphics capabilities (the CM-5 had none) for under $500,000 using commodity technologies. In addition, the 2001 system would have over 100 times as much storage capacity.

During that period, the cost-to-performance ratio for processing improved by a factor of 100. At the same time, networking, storage, display, and memory technologies improved by factors of 50 to 500. This extraordinary rate of improvement represents both opportunity and challenge, and the research and development (R&D) thrusts in MICS are aimed squarely at exploiting these improvements to provide better tools for energy researchers.

Finally, DOE's science research is increasingly done in teams of dozens to hundreds of participants. This process creates a critical need for technology that can support collaboration across time, distance, and the diversity of technical capabilities available to team members. A careful integration of research, development, and testing *at scale* is essential to provide useful tools to these large, distributed teams. This approach is particularly important because the components of the system improve at different rates, requiring a continual "rebalancing" of the overall system.

In addition to, and in conjunction with, operating advanced computing and communications facilities, MICS supports a set of interrelated research initiatives in the following areas.

Applied mathematics, computer science, and networking research involves creating computational methods, libraries, and techniques that fully exploit new computing, storage, display, and networking technologies. Work in these areas during the past decade was essential to enabling large-scale applications on parallel architectures. Current work is critical to the use of distributed systems, including not only commodity clusters, but also much larger (and perhaps even more promising) systems interconnected by ESnet and the Internet at large.

Enabling DOE science applications is the goal of several projects. In these projects, the results of fundamental research in applied mathematics and computer science are being translated into an integrated set of tools for application developers. These tools can be used by scientists in various disciplines to develop high-performance scientific applications (e.g., to simulate the behavior of materials) that will have a useful life spanning many generations of computer hardware. These tools will include capabilities for representing complex geometries, solving diverse numerical equations, simplifying multi-language parallel execution, evaluating and enhancing code performance, and dynamically steering calculations during execution.

Testbeds, collaboratories, and pilot projects develop the results of fundamental research in computer science and networking into an integrated set of tools to enable scientists to remotely access and control facilities and share data in real time. These activities include the development of robust "middleware" to support various kinds of grids (e.g., for collaboration environments, remote visualization and analysis of very large data sets, or creating large-scale distributed computational systems). To accomplish these goals, the following issues are being investigated:

- Definition and demonstration of a general and modular security architecture that can protect open network applications such as control of experimental devices.
- Development of a modular electronic notebook prototype that can be used in various desktop computer environments to enable the sharing of scientific results, data from scientific instruments, and design of scientific procedures.
- Development of tools to manage distributed collaborations. An important need is tools for managing multi-point videoconferences that range from the informal mode currently used, in which whoever is speaking has the floor, to more formal meetings, in which a meeting leader controls who has the floor.
- Development of advanced techniques for managing and returning to the electronic record of the collaboration.
- Exploration of techniques (e.g., virtual reality) to enable large groups to work together effectively at a distance.

Connectivity

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

NATIONAL LABORATORIES

Ames National Laboratory
Argonne National Laboratory
Brookhaven National Laboratory
Thomas Jefferson National Accelerator Laboratory
Los Alamos National Laboratory
Lawrence Berkeley National Laboratory
Lawrence Livermore National Laboratory
Oak Ridge National Laboratory
Pacific Northwest National Laboratory
Stanford Linear Accelerator Center

UNIVERSITIES

University of Arizona
University of Auburn
University of California, Berkeley
California Institute of Technology
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 Western Washington
Washington University
University of Wisconsin

INTERNATIONAL

Austria
Australia
Denmark
Japan
Norway
Poland
Sweden
United Kingdom

Facilities Requirements

Argonne Mathematics and Computer Science Division

The Mathematics and Computer Science (MCS) Division at Argonne National Laboratory provides both experimental research infrastructure and an integrated set of research activities that are organized into the Distributed Systems Laboratory, the Laboratory for Advanced Numerical Software, and the Futures Laboratory. MCS also supports other advanced facilities at the intersection of these three laboratories:

- Chiba City, a 512-processor commodity cluster for research concerning open-source scalability, applications performance, and management architecture.
- A set of advanced display, imaging, audio, visualization, and virtual environment facilities.
- I-WIRE (Illinois Wired/Wireless Infrastructure for Research and Education), a planned multi-pair, dark-fiber wide-area network, funded by the State of Illinois, to connect Argonne with the Urbana and Chicago campuses of the University of Illinois, the University of Chicago, and the STAR-TAP international transit point in Chicago.

MCS participates in collaborative projects in climate modeling, astrophysics, genomics, high energy physics, and other fields. These projects require access to high-speed networks for rapid movement of simulation and experimental data (a minimum of OC-12 today [622 Mbps] and moving to OC-48 [2.5 Gbps] and beyond by 2002). These groups will also need network engineering assistance to interconnect fast network caches and analysis engines at these line speeds.

Distributed Systems Laboratory (DSL). This laboratory, which is the home of the Globus multi-institutional computing effort, focuses on the tremendous potential of distributed supercomputing systems that couple parallel computers, high-performance workstations, large databases, virtual reality devices, and other resources connected by local or wide-area networks. DSL is dedicated to making such systems both usable and broadly accessible for computational science and engineering. To this end, DSL supports a broad range of research and development efforts, including high-level programming libraries and languages for parallel and distributed computing, as well as strategies for application-level management of advanced networking technologies (e.g., quality of service, security and scheduling technologies, and high-performance network protocols). DSL research on advanced middleware services requires access to experimental networks that support technologies such as differentiated services and MPLS and offer integrated instrumentation for monitoring network performance.

Futures Lab. Argonne's Futures Lab is developing prototype collaboration and visualization environments called "ActiveSpaces." ActiveSpaces are the workspaces of the future, places that combine existing workspace infrastructure with high-tech information technology. The goal is to construct a workspace that enhances the work experience, enables the user to be more productive, and does not intimidate the user. The construction of ActiveSpaces is a cross-cutting research agenda that combines research in display technology, collaboration environments, networking, and many other areas with the goal of creating a seamless environment.

Because of the growing number of research activities that require people to work in loosely coupled collections of small groups, the Futures Lab is developing working environments that enable groups to visually and interactively investigate large scientific data sets by using large-format and immersive visualization technologies in the context of shared collaborative spaces. These shared collaborative spaces have several attributes that distinguish them from current desktop-oriented, IP-based teleconferencing systems and traditional low-bandwidth videoconferencing. In particular, these spaces create the illusion of being in a shared workspace that is permanently connected to a set of other workspaces, via multiple cameras, large-format displays, and full-duplex "ambient audio" that allows participants to converse as if they were in the same work room.

Laboratory for Advanced Numerical Software (LANS). This laboratory conducts basic and applied research leading to new techniques and tools and new algorithms and software that fully exploit high-performance computers. This new technology will lay the groundwork for new scientific insights and approaches for solving challenging problems in science and engineering. LANS focuses on (1) developing analytical and numerical methods for solving complex physical systems and (2) building parallel programming tools that eliminate the barriers in developing applications for advanced computer systems. LANS software efforts include numerical software packages, parallel programming software, and program development software. Complementing this work in computer science and software development is research in applied mathematics and numerical analysis. Here the objectives are twofold: (1) explaining the underlying theory needed to understand complex phenomena and (2) analyzing the algebraic equations and mathematical models to ensure that they satisfy desirable properties such as accuracy, robustness, convergence, and efficiency.

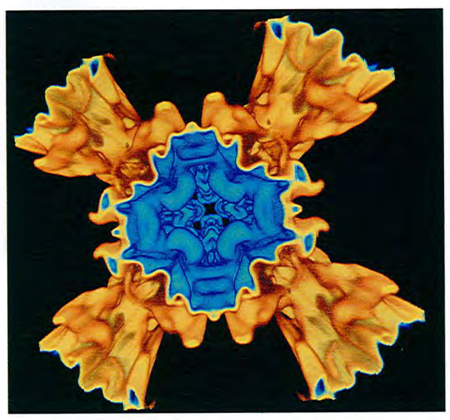

Volume rendering of the density variable from a simulation of a Rayleigh-Taylor instability performed by code developed at the University of Chicago's Center for Astrophysical Thermonuclear Flashes

National Energy Research Scientific Computing Center

The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory *(http://www.nersc.gov/)* is the flagship high-end computing facility for DOE's Office of Science. The driving force for current and future networking requirements at NERSC has shifted from computation to data storage and access. In the past, increases in archival storage needs were comparable to increases in computational capability, because the amount of data generated by computer simulations was usually limited by the available computational technology. Today this is no longer the case. Increasingly massive sets of experimental data are being generated by new technologies in fields such as genomics, climate research, high energy physics, and astrophysics. Computer centers are being called on to move, store, and analyze these data sets, and adequate network bandwidth is critical to meeting these needs.

The users of NERSC's largest High Performance Storage System (HPSS) have already commented that more network bandwidth is needed. In the FY 1999 NERSC user survey, "performance" and "response time" were the lowest-ranked categories for HPSS, with scores of 5.9 and 5.68, respectively (5 means "somewhat satisfied" on the scale of 1 to 7). Eleven of 38 users who commented on HPSS indicated a need for more bandwidth; specifically, seven cited a need to improve the transfer rate, and four said command response was too slow *(http://hpcf.nersc.gov/about/survey/99/hardware.html).*

NERSC's rule of thumb for the immediate future is that bandwidth must be at least doubled every year to meet the growing demands of researchers. This is a conservative estimate: for the commercial Internet, Cisco Systems projects a quadrupling of backbone requirements every year *(http://www.cisco.com/warp/public/cc/pd/rt/12000/prodlit/corbr_ov.htm).*

Thus, NERSC planned to double its bandwidth twice within calendar year 2000. At the main Berkeley Lab campus, NERSC has an OC-3 (155 Mbps) connection to ESnet. When NERSC moved its computing and storage systems to its new Oakland Scientific Facility in the fall of 2000, the connection was at 2x OC-3, and it will be upgraded to OC-12 (622 Mbps) in early 2001. To maintain its scientific and technological leadership, NERSC expects to need OC-48 (2.5 Gbps) bandwidth by 2002. Gigabit bandwidth is currently very expensive, but without it, the volume of scientific data will outstrip researchers' ability to analyze it.

Probe Distributed Storage Testbed

Turbulence structure downstream of an arterio-venous graft at a Reynolds number of 1820, visualized using the λ_2 method of Jeong and Hussain

The Probe Distributed Storage Testbed, a collaboration between Oak Ridge National Laboratory and NERSC, is representative of several major uses of ESnet:

- Developing and applying testing techniques to storage and networking hardware and software.
- Testing the use of ESnet 3 for bulk, high-bandwidth transfers over wide-area networks.
- Developing file transfer software to expedite transfer of files between HPSS installations.
- Developing a simulation model of local and wide-area storage/network operations.

Network and storage equipment, techniques, and software have evolved as rapidly as computers. Similarly, user demands for storage capacity, transfer rates, and high-bandwidth network access have grown rapidly. Storage and network systems have become extremely large and complex, often comprising thousands of hardware items.

This complexity, together with the rapid evolution of all elements, makes it extremely difficult to implement a cost-effective and high-performance storage system with high-bandwidth access for all customers. To address the complexity issue, we have initiated a project to develop a simulation model of HPSS and, in particular, its use in challenging storage applications. This model will incorporate local and wide-area network features so transfers, for instance, to local and distant visualization sites can be studied and optimized.

Projects already completed or underway include testing and tuning the performance of the HPSS with gigabit Ethernet transfers, testing transfers between storage and supercomputers through a Gigabyte System Network (GSN, also known as HIPPI-6400), testing and tuning of the ESnet link between Oak Ridge National

Laboratory and NERSC, and collaborative activities (such as beta tests) with storage and network equipment vendors. Future tests will investigate scheduled transfer, 10-gigabit Ethernet networks, InfiniBand (Intel), and storage over IP as those technologies become available.

Research Requirements

MICS is expanding its network research, as outlined in Chapter 2 of this plan. This suite of network research initiatives both depends upon and drives other MICS research in the areas outlined below.

Applied Mathematics, Computer Science, and Networking

Advanced Computational Testing and Simulation (ACTS) Toolkit. The ACTS Toolkit is a set of DOE-developed software tools that make it easier for programmers to write high-performance scientific applications for parallel computers. Parallel software is inherently more complex than serial software and significantly more expensive to implement. In the past, many smaller projects have been unable to access the benefits of parallel computing because they lack the resources to convert their computational codes from serial to parallel implementations. Moreover, as the problems being solved become more complex, the challenges in building parallel programs become even greater. The goal of the ACTS Toolkit is to alleviate this situation. Moreover, because of the improved price-to-performance ratio of both computing and networking resources, it has become possible to create both large-scale PC clusters and wide-area coupled clusters. In this sense, the line between "parallel" and "distributed" computers continues to blur. With the widespread and growing use of these technologies to build "computational grids," the ACTS Toolkit and other applied mathematics and computer science efforts supported by MICS will become fundamentally important not only for traditional computers but also for distributed systems.

Collaborative technologies. MICS also supports a wide range of research and development in collaborative technologies. To make existing resources truly effective, scientists and engineers must be able to interact as if they were physically collocated—sharing data, high-performance computing systems, and instrumentation independent of location. MICS has supported a variety of activities via the DOE2000 collaboratory research program. MICS collaboratory research and development efforts include interoperability frameworks, electronic logbooks, collaborative session management, shared virtual spaces, scalable security architecture, and network quality of service.

High-bandwidth networks are obviously critical to collaboratory research because of its heavy emphasis on multimedia. However, bounded delay and other quality-of-service metrics are essential not only for collaboration technologies but also for distributed computing systems. Both depend on predictable, low-latency messaging—for efficiency in the case of distributed computing and for ease of human interaction in the case of collaborative systems.

Enabling DOE Science Applications

MICS research and development in enabling applications is at the vanguard of an ongoing revolution in the way that computing and networks, or "grids," are used to support DOE science. In this work, new computational techniques and tools, such as those in the ACTS Toolkit discussed above, are applied to basic research problems to test the usefulness of current advanced computing research. Project teams include both computer scientists and researchers in the discipline being studied. This team structure has proved to be an effective way to transfer the results of this research to the scientific disciplines and to define promising areas for future research. Some teams investigate theoretical and modeling topics, such as

- Climate simulations.
- Combustion modeling.
- Fundamental structure and properties of magnetic materials.
- Advanced tools to understand the chemistry of actinides.

Other partnerships involve experimental disciplines such as

- High energy and nuclear physics.
- Human genomics.
- Crystallography.

These projects are intended to help scientists in these varied disciplines manage and analyze the petabyte-scale data sets produced by their experiments and simulations (see page 101). This work has evolved from the Grand Challenge projects that were initiated as part of DOE's component of the Federal High Performance Computing and Communications program, which started in FY 1991.

Testbeds, Collaboratories, and Pilot Projects

MICS scientists participate in various testbed, collaboratory, and pilot projects in partnership with other DOE programs. The goal of these projects is to test, validate, and apply the basic research tools and techniques created through the efforts described above. Two examples are the Materials Microcharacterization Collaboratory and the Diesel Combustion Collaboratory, both of which are partnerships between the Office of Science and the Office of Energy Efficiency and Renewable Energy in DOE. Both aim to provide collaboration services and remote access to experimental and computing facilities for teams of researchers from national laboratories, universities, and industries (see the Basic Energy Sciences section of Chapter 3). Originally a part of the DOE2000 initiative, these partnerships with the scientific disciplines and the technology programs within DOE enable the MICS Division to be particularly successful in bringing the results of computing research to bear on important scientific problems. In addition, they provide important feedback to MICS researchers on which problems most warrant further study.

The CorridorOne Project: Computational steering for data-intensive applications

CorridorOne combines the forces of six leading-edge laboratory and university groups working in the areas of visualization, distributed computing, and high-performance networking. The CorridorOne team is attempting to develop the most advanced integrated distance visualization environment for large-scale scientific visualization. The technology will be demonstrated in applications that are critical to DOE's science programs and Accelerated Strategic Computing Initiative (ASCI).

***Strategy.** The strategy for this project builds on the success of the teams that produced the I-WAY, Computational Grids, the award-winning GUSTO testbed, and CAVE display technology. This expertise will be expanded by adding the work of teams that have developed the fastest parallel visualization systems and the most widely used networking infrastructure for multicast and distributed media.*

***Key Activities.** The CorridorOne team will prototype and demonstrate a six-way multi-point distance visualization corridor built on top of a state-of-the-art Grid fabric. The corridor will be used to experiment with and demonstrate a variety of high-performance remote visualization technologies on challenging applications. CorridorOne is working directly with a number of ASCI applications. Large-scale data sets, such as modeling data from climate, fusion, and combustion research, are driving the development of CorridorOne.*

The team is working closely with groups supported by DOE, NSF, and NASA to deploy Grid services at the CorridorOne sites. CorridorOne is also developing network technologies, such as differentiated network services, multi-domain authentication and resource brokering services, adaptive network APIs (application programming interfaces), and high-performance transport protocols, to support existing large-scale visualization tasks.

CorridorOne Focus Areas

Grid middleware and advanced networking—Accelerating deployment of Grid middleware services to the CorridorOne sites and developing Grid infrastructure and middleware technology to provide requirements and feedback on capabilities needed to support distance visualization.

Distributed visualization and data manipulation—Exploring technologies aimed at both reducing the amount of data that needs to be transmitted across the network to accomplish some visualization task (e.g., raw data, geometry, and images) and reducing the latency for interaction and navigation through large data sets (e.g., progressive refinement, multi-resolution, and feature-to-feature navigation).

Distributed collaboration and display environments—Enabling researchers from multiple locations to interpret and understand the results of large-scale calculations.

Systems architecture, software frameworks, and tool integration—Developing a high-level systems architecture that can support the divergent needs of both a stable software base for experimentation and applications development, as well as a more fluid environment for exploring new concepts.

Applications liaison, experimental design, and evaluation—Providing test data and users (from six application groups) that span domains of interest to DOE.

CorridorOne Project Participants
Argonne National Laboratory
Lawrence Berkeley National Laboratory
Los Alamos National Laboratory
Princeton University
University of Illinois at Chicago
University of Utah


Data courtesy of M. Wheeler and M. Peszynska, University of Texas

Image of 10 isosurfaces of oil concentration in production oil wells (yellow) and injection wells (blue) computed on a commodity cluster using the Visualization Toolkit and associated extensions

BOOMERANG: Astrophysical observations show the future of scientific computing

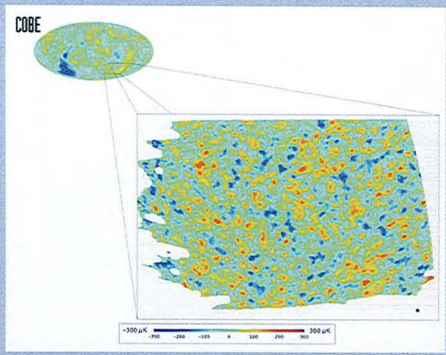

Sound waves in the embryonic universe are revealed for the first time in this image captured by the BOOMERANG telescope

By using high-performance computing to analyze data from a balloon flight in the Antarctic, scientists gained new insight into the geometry of our universe—and a glimpse at the future of scientific computing. In 1999, as part of an international research project, the BOOMERANG balloon mission measured cosmic microwave background data. It generated massive data sets: 50 million observations for each of 16 channels at four frequencies. The National Energy Research Scientific Computing Center (NERSC) and ESnet provided essential tools for analyzing the data.

Using these tools, the BOOMERANG team was able to make the most detailed map ever seen of the temperature fluctuations in this background radiation *(http://www.physics.ucsb.edu/~boomerang/)*. From such a map, the researchers derived a "power spectrum," a curve that registers the strength of the fluctuations on different angular scales. It also yields information on such characteristics of the universe as its geometry and how much matter and energy it contains. The power spectrum allowed researchers to conclude with a high degree of confidence that the geometry of the universe is flat.

BOOMERANG maps the cosmic microwave background using a balloon-borne telescope that circumnavigates Antarctica

These calculations required 50,000 hours of processor time on NERSC's Cray T3E supercomputer, a job that would have taken six years to complete on the most powerful desktop computer. Future missions will produce data sets tens to hundreds of times bigger than those from the BOOMERANG project. Two examples are NASA's Magnetic Anisotropy Probe (MAP), to be launched in 2001, and the European Space Agency's PLANCK, to be launched in 2007. These massive data sets will require new computational strategies and will place heavy demands on network capabilities—challenges that will need to be addressed by NERSC and ESnet.

3 Core Program Requirements

Defense Programs

Program Description

Mission

The mission of DOE's Defense Programs (DP) is to ensure the safety, reliability, and performance of nuclear weapons without underground nuclear testing *(http://www.dp.doe.gov/dp_web/)*. 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 begun developing technologies in support of its national security mission under requirements of the Comprehensive Test Ban Treaty (CTBT). This mission 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 based on advanced modeling and simulation methods, complemented by aboveground experiments and validated against historical test data.

Computing

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 in allowing stockpile stewardship to shift promptly from test-based to science-based assessment methods. It will result in advanced computational capabilities essential for successful stewardship. For example, new three-dimensional applications using more detailed physics are being created. The data produced by these models will vastly exceed today's data sets in both quantity and complexity and will require innovative approaches to data analysis and data assimilation.

The ASCI program is executed at Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratories 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 facilities are located at Oak Ridge National Laboratory, Pantex, Allied Signal, and Savannah River.

ASCI employs five strategies to achieve its objectives:

- Create seamless management: one program, three laboratories.
- Focus on advanced applications development.
- Focus on the high end of computing.
- Create problem-solving environments.
- Encourage strategic alliances and collaborations.

Program Requirements

Capacity

Through most of FY 2000, ESnet has provided the single connection for unclassified communications between both the laboratories and the plants. This OC-3 (OC-12 at Los Alamos) network between the laboratories has enabled both the necessary development work related to accomplishing remote supercomputing as well as actual production work. Concurrently, this network has provided the means to move manufacturing design information between the laboratories and the plants.

In FY 1999 and 2000, design was completed for a parallel networking architecture that will be used to provide secure high-speed (2.5 Gbps to 100 Gbps) machine-to-machine connections between the three laboratories. In June 2000, contracts were awarded to provide the necessary wide-area network component of this architecture. The intent is that the current ESnet service will continue to be used for all of the DP traffic that is not in the category of machine-to-machine (computing platforms, storage, and visualization) data transfer. The ESnet bandwidth requirement between all of the laboratories will increase to OC-12 (622 Mbps) in 2001.

Connectivity

DP requires reliable high-speed open networks (in addition to separate secure classified networks) that interconnect the three laboratories, five University Alliance Centers, and four manufacturing plants. These requirements range from providing normal Internet services (collaboration tools, e-mail, etc.) among all the sites to moving terabytes of data between the three laboratories in connection with running large weapon simulation codes remotely on the ASCI computational platforms.

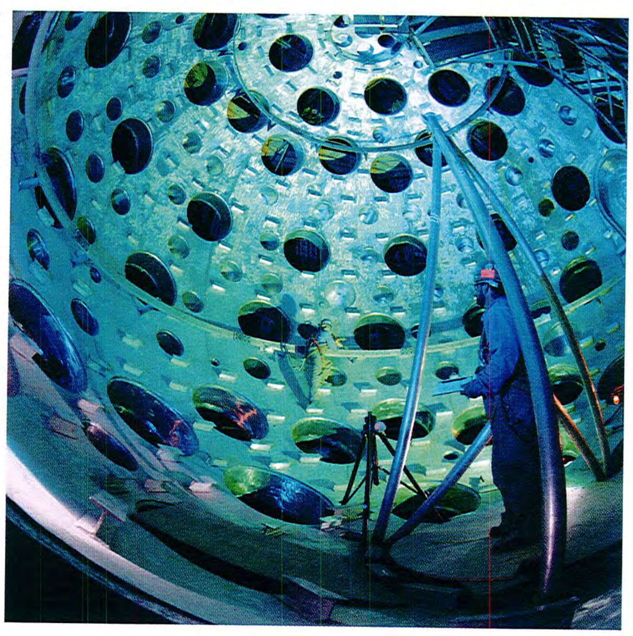

Photo courtesy of J. McBride and B. Quintard, Lawrence Livermore National Laboratory

Experiments at the National Ignition Facility (NIF) will rely on supercomputing supported by ESnet. Photo shows the interior of the target chamber, which weighs one million pounds and measures 30 feet in diameter.

4 Other DOE Program Requirements

Chief Information Officer

The Office of the DOE Chief Information Officer (CIO) develops the overall information technology policy under which DOE operates *(http://cio.doe.gov/)*. The Office of the CIO also coordinates planning for major information technology investments by the Department. As a representative of the Department as a whole, the CIO seeks to develop infrastructure, processes, and services that will help a wide range of DOE programs meet their mission objectives and that will support DOE energy strategies and the national interest.

International Nuclear Safety Program

DOE manages a comprehensive collaborative effort to improve nuclear safety at Soviet-designed nuclear power plants in nine partnering countries. In addition to these nine host countries, the G-7 nations and many international organizations collaborate on this effort.

The principal goals of this International Nuclear Safety Program (INSP) are to improve safety and reduce risks by

- Improving physical operating conditions.
- Installing safety equipment.
- Developing improved safety procedures.
- Establishing regional centers for training reactor personnel.
- Installing simulators for training control room operators.
- Conducting in-depth safety assessments.
- Developing institutional and regulatory frameworks.
- Addressing the extraordinary problems at Chornobyl.

For more than a decade, the world community has worked to assess the long-term effects of the 1986 disaster at Chornobyl, including its economic, environmental, and health impacts. To focus these efforts, Ukrainian President Leonid Kuchma established the Chornobyl Center for Nuclear Safety, Radioactive Waste, and Radioecology in 1996. The Center's primary technical branch, the Slavutych Laboratory for International Research and Technology, is 40 kilometers northeast of the Chornobyl Nuclear Power Plant, in the city of Slavutych.

To facilitate communications, DOE has worked with Ukraine to establish a reliable, high-speed satellite link that enables transmission and reception of voice, facsimile, Internet, and videoconferencing data. The satellite link terminates at Pacific Northwest National Laboratory, where it interfaces with ESnet. ESnet provides the connectivity to the DOE laboratories and other sites in support of INSP initiatives.

Office of Counterintelligence

The DOE Office of Counterintelligence (DOE-CN) maintains equipment and systems specifically designed to support its counterintelligence mission. This office works with the Federal Bureau of Investigation (FBI) and contributes cyber threat information to the National Infrastructure Protection Center (NIPC) so these organizations can better protect the critical national infrastructures as defined in Presidential Decision Directive (PDD) 63.

In this work, ESnet is used to send information from DOE sites to analysis facilities and then to the NIPC at the FBI. As a result, ESnet will play an important role in providing the means to better protect critical national infrastructures from cyber threats.

Photo courtesy of National Aeronautics & Space Administration

Scientists now depend on reliable and secure global communications in their daily work

ACTS	Advanced Computational Testing and Simulation [Toolkit]
AGS	Alternating Gradient Synchrotron, BNL
ALB	Albuquerque
ALS	Advanced Light Source
AMES	Ames Laboratory
ANL	Argonne National Laboratory
APS	Advanced Photon Source
ARM	Atmospheric Radiation Measurement [Program]
ASCI	Accelerated Strategic Computing Initiative
ASDEX-U	Axially Symmetric Divertor Experiment–Upgrade
ASIG	Allied Signal
ATL	Atlanta
ATLAS	Argonne Tandem-Linear Accelerator System
	A Toroidal LHC ApparatuS, CERN
ATM	Asynchronous transfer mode
BaBar	B/B-bar detector, SLAC
BEPC	Beijing Electron-Positron Collider
BER	Office of Biological and Environmental Research, DOE
BES	Office of Basic Energy Sciences, DOE
BITnet	"Because It's Time" Network
BNL	Brookhaven National Laboratory
BOOMERANG	Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics
bps	Bits per second
BRAHMS	Experiment at RHIC
CAVE	A virtual reality environment
CCPP	Climate Change Prediction Program
CCTI	Climate Change Technology Initiative
CEA	Commissariat à l'Energie Atomique
CERN	European Organization for Nuclear Research
CESR	Cornell Electron Storage Ring
CHI	Chicago
CHI-NAP	Chicago Network Access Point
CIO	Chief Information Officer
CLAS	Detector at Jefferson Lab
CMM	Center for Microanalysis of Materials
C-Mod	Tokamak at MIT
CMSN	Computational Materials Sciences Network
CRADA	Cooperative Research and Development Agreement
CRF	Combustion Research Facility
CRPP	Centre de Recherches en Physique des Plasmas
CTBT	Comprehensive Test Ban Treaty
DARPA	Defense Advanced Research Projects Agency
DCC	Diesel Combustion Collaboratory

Acronyms and Abbreviations

DCS	Digital collaboration services
DCWG	Distributed Computing Working Group
DESY	German Electron Synchrotron Laboratory
DOE	U.S. Department of Energy
DP	Office of Defense Programs, DOE
DRFC	Département de Recherches sur la Fusion Contrôlée
DSL	Distributed systems laboratory
DWDM	Dense wave division multiplex
EIA	Energy Information Agency
EMC	Electron Microscopy Center
EMSL	William R. Wiley Environmental Molecular Sciences Laboratory
ENEA	Ente per le Nuove Technologie, l'Energie e l'Ambient
EOS	Earth Observing Satellite
EPFL	Ecole Polytechnique Fédérale de Lausanne
EPSCoR	Experimental Program to Stimulate Competitive Research
ER	Office of Energy Research, DOE (now Office of Science)
ESCC	ESnet Coordinating Committee
ESD	Environmental Sciences Division, DOE
ESnet	Energy Sciences Network
ESRSC	ESnet Research Support Committee
ESSC	ESnet Steering Committee
FBI	Federal Bureau of Investigation
FDDI	Fiber distributed data interface
Fermilab	Fermi National Accelerator Laboratory
FES	Fusion Energy Sciences
FIX-W	Federal Interagency eXchange-West
FNAL	Fermi National Accelerator Laboratory (Fermilab)
FY	Fiscal year
GA	General Atomics
Gbps	Gigabits per second
GCM	General circulation model
GeV	Giga electron volts
giga-	10^9 or 1 billion
GIOD	Globally Interconnected Object Database
GOES	Geostationary Operational Environmental Satellite
GSN	Gigabyte System Network
GTN	Germantown
H.323	Protocol for IP-based videoconferencing
HEP	High Energy Physics
HEPnet	High Energy Physics Network
HERA	Hadron Elektron Ring Anlage, DESY
HFBR	High Flux Beam Reactor
HFIR	High Flux Isotope Reactor

HGPI	Human Genome Project Information
HIPPI	High-performance parallel interface
HPSS	High-performance storage system
IEEE	Institute of Electrical and Electronics Engineers
IEPM	Internet end-to-end performance monitoring
IETF	Internet Engineering Task Force
IFE	inertial fusion energy
IHEP	Institute for High Energy Physics
INEEL	Idaho National Engineering and Environmental Laboratory
INSP	International Nuclear Safety Program
IP	Internet protocol
IP/SONET	Direct Internet protocol over synchronous optical network
IPNS	Intense Pulsed Neutron Source
IPP	Institute for Plasma Physics
IPv4(6)	Internet protocol version 4 (or 6)
I-WIRE	Illinois Wired/Wireless Infrastructure for Research and Education
JAERI	Japanese Atomic Energy Research Institute
Jefferson Lab	Thomas Jefferson National Accelerator Facility
JET	Joint European Torus
JGI	Joint Genome Institute
JINR	Joint Institute for Nuclear Research
JT	Japanese Tokamak
kbps	Kilobits per second
KEK	Japanese National Laboratory for High Energy Physics
kilo-	10^3 or 1,000
KSTAR	Korean Superconducting Tokamak Advanced Research
LAN	Local-area network
LANL	Los Alamos National Laboratory
LANS	Laboratory for Advanced Numerical Software
LANSCE	Los Alamos Neutron Science Center
LBNL	Lawrence Berkeley National Laboratory
LEP	Large Electron-Positron Collider, CERN
LHC	Large Hadron Collider, CERN
LHD	Large Helical Detector
LLNL	Lawrence Livermore National Laboratory
MAE-E	Metropolitan Access Exchange-East
MAE-W	Metropolitan Access Exchange-West
MAP	Magnetic Anisotropy Probe
MAST	MegaAmp Spherical Tokamak
MB	Megabytes
Mbps	Megabits per second
mega-	10^6 or 1 million
MEM	Memphis

Acronyms and Abbreviations

MFE	Magnetic fusion energy
MFEnet	Magnetic Fusion Energy Network
MHD	Magnetohydrodynamics
MICS	Mathematical, Information, and Computational Sciences Division, DOE
MIT	Massachusetts Institute of Technology
MMC	Materials Microcharacterization Collaboratory
MPC	Materials Preparation Center
MPLS	Multi-protocol label switching
NABIR	Natural and Accelerated Bioremediation Research
NAP	Network access point
NASA	National Aeronautics and Space Administration
NCEM	National Center for Electron Microscopy
NCSX	National Compact Stellarator Experiment
NERSC	National Energy Research Scientific Computing Center, LBNL (formerly National Energy Research Supercomputer Center at LLNL)
NIFS	National Institute for Fusion Studies
NIPC	National Infrastructure Protection Center
NOAA	National Oceanic and Atmospheric Administration
NSF	National Science Foundation
NSLS	National Synchrotron Light Source
NSTX	National Spherical Torus Experiment
NTCC	National Transport Code Collaboration
NYC	New York City
NY-NAP	New York Network Access Point
OASCR	Office of Advanced Scientific Computing Research
OC-12	622 Mbps
OC-3	155 Mbps
OC-48	2.5 Gbps
ORNL	Oak Ridge National Laboratory
PB-NAP	Pacific Bell Network Access Point
PEP-II	Positron-Electron Project II, SLAC
peta-	10^{15} or 1,000 trillion
PHENIX	Experiment at RHIC
PHOBOS	Experiment at RHIC
PingER	Ping end-to-end reporting
PKI	Public key infrastructure
PMT	Photomultiplier tube
PNNL	Pacific Northwest National Laboratory
POS	Packet over SONET
pp2pp	Experiment at RHIC
PPPL	Princeton Plasma Physics Laboratory
PSACI	Plasma Science Advanced Computing Initiative
QoS	Quality of service

R&D	Research and development
RCWG	Remote Conferencing Working Group
REDC	Radiochemical Engineering Development Center
RHIC	Relativistic Heavy Ion Collider, BNL
RIKEN	Institute of Physical and Chemical Research (Rikagaku Kenkyusyo)
SAIC	Science Applications International Corporation
SC	Office of Science, DOE
SciDAC	Scientific Discovery through Advanced Computing
SDSC	San Diego Supercomputing Center
SEA	Seattle
ShaRE	Shared Research Equipment Program
SLAC	Stanford Linear Accelerator Center
SMAC	Surface Modification and Characterization Research Center
SNL	Sandia National Laboratories
SNLA	Sandia National Laboratory (Albuquerque)
SNLL	Sandia National Laboratory (Livermore)
SNS	Spallation Neutron Source
SONET	Synchronous optical network
SRS	Savannah River Site
SSRL	Stanford Synchrotron Radiation Laboratory
STAR	Experiment at RHIC
Super-K	Super Kamiokande
T1	1.5 Mbps
T3	45 Mbps
Tbps	Terabits per second
TCP	Transmission control protocol
tera-	10^{12} or 1 trillion
TeV	Trillion electron volts
UAV	Unmanned aerospace vehicle
UCLA	University of California at Los Angeles
UKAEA	United Kingdom Atomic Energy Authority
vBNS	Very high speed Backbone Network Service
VDG	Virtual data grid
VPN	Virtual private network
WAN	Wide-area network
YUCCA-MT	Yucca Mountain

A History of Excellence

The impetus for forming ESnet developed in the mid-1980s, when both the Fusion Energy (FE) and High Energy Physics (HEP) programs in DOE 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 after a dedicated FE supercomputer center opened in 1974 at Lawrence Livermore National Laboratory (LLNL). HEP researchers had similarly constructed a home-grown network (HEPnet), starting in the late 1970s, with microwave links between the Stanford Linear Accelerator Center (SLAC) and Lawrence Berkeley National Laboratory (LBNL), followed shortly by a satellite link between SLAC and Argonne National Laboratory. By the mid-1980s, the HEP program had developed an extensive network of leased lines (mostly operating at 9600 bits per second [bps]) that interconnected the main particle accelerator laboratories with many other sites.

The early 1980s also saw other DOE Office of Energy Research (ER) programs joining established computer networks. Many national laboratory and 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 whose research was concentrated at remote laboratories found it necessary to lease direct connections to mainframe computers at those laboratories.

In fiscal year (FY) 1985, Dr. Alvin Trivelpiece, then Director of the Office of Energy Research, charged ER's Scientific Computing staff (now the Mathematical, Information, and Computational Sciences Division) 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 capabilities 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.

As a result of these findings, the MFEnet and HEPnet initiatives were combined and broadened into what would become ESnet, a single general-purpose scientific network for the ER community. The ESnet Steering Committee (ESSC) was formed to be a source of guidance on requirements and general strategic oversight. The installation, coordination, and day-to-day operation of ESnet became the responsibility of the National Magnetic Fusion Energy Computer Center, which was renamed the National Energy Research Supercomputer Center to reflect its expanded role of providing supercomputer access and network services to a wider community.

The ESSC held its first meetings in late 1986, and ESnet began providing ER-wide networking services in January 1988. Starting with X.25 backbone lines at speeds of 56 kilobits per second (kbps) and 256 kbps, ESnet grew to T1 backbone lines (1.5 Mbps) in 1989, and T3 (45 Mbps) in 1993. In February 1995, ESnet was the first major network to convert to a new technology—asynchronous transfer mode (ATM)—and today more than 30 sites are directly connected at speeds as high as OC-3 (155 Mbps) and OC-12 (622 Mpbs).

During the late 1980s and the early 1990s, ESnet also began providing international connectivity in support of ER program activities. This connectivity was shared and

Networking accomplishments

DOE has a long and distinguished history of supporting network R&D to meet its own science mission. Equally important, the results of such research have contributed significantly to the success of the global Internet. The Department's researchers have made seminal contributions to key Internet technologies, including the following:

- *TCP (transmission control protocol) congestion management.*
- *"Random early detection" router algorithms that are universally used today.*
- *Class-based queuing (the foundation of the IP quality of service found in most commercial IP routers).*
- *IP multicast-based video and audio collaboration tools.*

More recently, continuing this tradition of advancing the capabilities that allow DOE scientists to perform new science, ESnet

- *Has created the DOE Science Grid Network, which is the first quality-of-service testbed based on multi-domain differentiated services.*
- *Leads the efforts by the Internet Engineering Task Force (IETF) to transition from the Internet protocol version 4 (IPv4) addressing architecture to the next generation, IPv6, and was awarded the first permanent production IPv6.*
- *Has established and maintains a DOE-wide digital collaboration service, as well as the first DOE-wide virtual private network.*
- *Has served as a catalyst to major communication carriers to provide asynchronous transfer mode (ATM) technology nationwide.*
- *Fosters and supports international peering with more than 15 countries.*
- *Provides peering points for more than 150 Internet service providers to connect with DOE resources in the United States.*
- *Facilitates connections between universities and DOE by peering with the National Science Foundation's very high speed Backbone Network Service (vBNS) and Abilene, a high-performance research network for universities (http://www.ucaid.edu/abilene/).*
- *Supports an annual growth in traffic of nearly 100% per year. Total traffic over ESnet for calendar year 2000 will exceed 400 terabytes, the equivalent of about 600,000 compact discs.*

The control room of the Large Helical Device (Toki, Japan), the largest fusion energy experiment of its type and an important focus of collaborative computing coordinated with the National Aeronautics and Space Administration (NASA), National Science Foundation (NSF), and Defense Advanced Research Projects Agency (DARPA). This approach established a framework for future interagency cooperation. In the same period, ESnet began connecting to regional NSF networks, thereby providing more ubiquitous network communications through which university researchers could make use of ER laboratories and facilities.

Today, as an integral part of the Internet, ESnet provides seamless, multi-protocol connectivity among a variety of scientific facilities and computing resources in support of collaborative research. The benefits that have accrued to the scientific community have only whetted researchers' appetite for transmitting ever-greater volumes of information at faster rates and using more sophisticated applications. The mission of ESnet is to satisfy these needs as fully as possible for DOE researchers.

ESnet Organization and Process

A successful service-oriented community must first establish user requirements and then effectively satisfy those requirements within operating, funding, and administrative constraints. For ESnet, these processes involve one of the most widespread collaborative efforts within DOE.

Assessments of current and future user requirements are obtained from DOE-supported principal investigators located at national laboratories and universities across the United States 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 (LBNL), 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 e-mail and videoconferencing support—requires a technically diverse and knowledgeable support staff. The complex interactions of these technical experts are overseen by an effective, efficient management and committee structure whose hallmarks are cooperation and synergy.

Mathematical, Information, and Computational Sciences Division

ESnet is sponsored by the Mathematical, Information, and Computational Sciences Division (MICS) within the Office of Advanced Scientific Computing Research (OASCR) of DOE's Office of Science (SC). 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 SC's supercomputer facilities and ESnet.

ESnet Steering Committee

History and Mission

The ESnet Steering Committee (ESSC) was originally formed in 1986 with representation from the Energy Research (now Office of Science) program offices. As ESnet has extended its service to other parts of DOE, the ESSC has expanded to ensure adequate breadth of program representation.

The ESSC is charged to

- Document, review, and prioritize network requirements for all SC 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 management personnel at LBNL.

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 were developed in the formative stages of the ESSC and are documented in the minutes of its meetings.

Definition and Validation of Requirements

The ESSC has codified the process and criteria by which priorities can be set for major network improvements, including major new network connections as well as upgrades of existing facilities. The 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 support the new or expanded network connectivity.

The ESSC has documented these criteria and 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.

Requirements generally originate from programs or from sites. Committee members identify requirements from the programmatic community they represent and present these requirements at ESSC meetings. The requirements are then reviewed by the whole committee, which takes into account the needs of the programs and the availability of resources. When worthy requirements cannot be accommodated immediately, the committee uses a prioritization process to determine the order in which requirements will be addressed.

Requirements that refer in general to one of the DOE laboratories or other ESnet sites, rather than to a specific program, are often identified through the ESnet Site Coordinating Committee, the ESSC subcommittee representing the sites. The site requirement is referred to ESSC members to determine connections with each of the programs. In the end, site requirements are related to programmatic requirements and reviewed on their programmatic merit.

The committee regularly hears reports on network research and other technology advances. Research on networking in DOE has a special relationship with ESnet, not only because of its programmatic proximity in the MICS program office, but also because advances made in the DOE program can often be deployed especially quickly, giving ESnet a technical advantage.

Planning Process for Future Needs

The activities of the ESSC are designed to ensure that future networking needs of DOE programs are met. The committee tracks each of its participating programs on a regular basis in order to be alert to changing needs. Documentation of future needs is provided in the program plan (this document) at intervals of roughly three years and in the ESnet strategic plan.

An important element of planning is the use of task forces and working groups to deal with specific issues as they arise. Recent issues addressed in this way include

- Advanced applications.
- University support.
- Network research testbeds.
- International connections.
- Publicity.
- Strategic plan.
- Program plan.
- Security.

In addition, the ESSC continually addresses general issues of ongoing significance for the future of DOE networking. Recently these issues have included

- Telecommuting.
- Videoconferencing and collaborative 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 SC community.
- The relationship of ESnet to national initiatives.
- The effects of related federal activities.

The ESSC also considers such strategic issues as

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

Standing Subcommittees

The ESSC has two standing subcommittees to assist with particularly important and continuing aspects of its mission.

ESnet Site Coordinating Committee (ESCC). The first standing subcommittee is the ESnet Site Coordinating Committee, which was founded in 1987 with the original purpose of providing coordination with the sites that needed to work technically and administratively with ESnet. Because the site representation included extensive technical expertise, the mandate of the subcommittee was quickly extended to include technical working groups and consultation with ESnet management. The new role included an emphasis on providing tests of emerging technical developments that were expected to benefit the operation of the network. The technical consultation function has developed into an extensive portfolio of working groups that constitute a major source of grassroots participation in ESnet. Such consultation has contributed materially to the success of ESnet.

Members of the ESCC are appointed by the individual ESnet site organizations. Current membership represents 22 major sites. The ESCC chairperson is appointed by MICS from among the members of that committee, with the advice and consent of both the Steering Committee and the ESCC. The ESCC serves as

- An advisory body to the Steering Committee, providing a forum for considering 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 SC programs that use or would like to use ESnet facilities.

The committee uses an approach based on working groups and task forces. Task forces tend to be groups formed to work on a specific topic for a relatively short duration, while working groups tend to have a longer life span. Thus, for example, the present authentication task force is assumed to have a finite life span, while the architecture working group should exist for some time. The authentication task force is working through both the policy and technology issues associated with inter-laboratory authentication and security domains. Similarly, in the past, task forces have been charged with migrating away from DECnet and studying the installation of ATM. When these projects were completed, the task forces disbanded.

ESnet Research Support Committee (ESRSC). The second standing subcommittee, formed in 2000, is the ESnet Research Support Committee, whose members are appointed by the Steering Committee. The ESRSC provides a forum in which network researchers, scientific users, and ESnet staff members can meet to coordinate the introduction of new network protocols and services into ESnet. The ESRSC is chartered as

- An advisory body to the Steering Committee that provides a guiding vision and advocacy for research activities that will advance DOE distributed computing and collaborative environments.
- An advisory body to the Steering Committee that identifies the program requirements and opportunities for testing expanded network capabilities.
- A forum for coordinating ESnet testbed activities and recommending the deployment of new research support services.

Program Reviews

The DOE MICS office, which sponsors ESnet, and the Steering Committee have placed considerable emphasis on regular reviews of the program. The Program Office and the Steering Committee have cooperated in holding formal reviews of ESnet at intervals of typically three years. The reviews have been conducted by well-known members of the computing and networking communities, as well as by active researchers from the DOE programs served by ESnet. Reviews have focused on different areas over the years. For example, reviewers have been asked to evaluate technical directions, project management, responsiveness to DOE missions, community organization and committee structure, and many other indicators of effectiveness. While all reviews have been quite positive about the accomplishments of ESnet, these reviews by outside experts have paid off with many valuable recommendations that have led to new directions or other changes in the ESnet program. Along with its unique committee structure, program reviews by outside experts have kept ESnet technically strong and able to anticipate the diverse needs of its users.

ESnet and the Larger Internet Community

Partnerships and Collaborations

ESnet continues to participate in the development of the global Internet. ESnet has historically been a contributor within that community, having been an early explorer of new technologies such as dual-protocol routers and ATM. By successfully pioneering the use of such emerging technologies in its production-quality network, ESnet helps advance the state of technology within the wider Internet community.

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

ESnet personnel participate in the Internet community forums, task forces, working groups, committees, and subcommittees in which standards, protocols, and acceptable practices are developed. For example, ESnet has been represented on the

- Internet Engineering Task Force.
- North American Network Operators Group.
- Joint Engineering Team.

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

- Advanced protocols.
- Internet security.
- Congestion analysis and avoidance.
- Public key infrastructure.
- IPv6 (Internet protocol version 6).
- Operations management.
- Directory services.
- Privacy.
- Quality of service (QoS).
- Application support.

In addition, many management and joint engineering opportunities have resulted in mutually beneficial arrangements with other parts of DOE, other government agencies, and the international research community. ESnet is widely recognized for the technical competency of its staff and its ability to find practical solutions to technical, financial, and political challenges.

Memoranda of Understanding

Through the budget process, ESnet funding is appropriated to support SC mission requirements. This funding is not used to support other major DOE programs or organizations unless otherwise provided for through memoranda of understanding. However, such agreements for coordinated and cooperative use of the network infrastructure are encouraged because they help DOE achieve larger goals, such as improving access to government information, minimizing overlap, and complementing other communications requirements.

International Collaboration

The programs of DOE's Office of Science (the largest users of ESnet) are strongly international. The international character appears in different ways, creating different requirements for ESnet. Researchers have always needed to communicate with peers in many other countries to share ideas and to compare results. Increasingly, however, the norm is active collaboration among researchers in two or more countries who work together directly on a particular research project. This extended collaboration may come about because the participants need to share a unique facility or other resource, because the researchers bring different knowledge and expertise to the joint enterprise, or because the project is simply too big for one institute or even one country to take on comfortably.

From a networking point of view, the collaborations that generate the greatest requirements are those focused on unique major facilities such as an accelerator, fusion test facility, or neutron source. Many projects vital to the DOE science mission are carried out by international teams either at a DOE facility, for example, the Advanced Photon Source or the Stanford Linear Accelerator Center, or at a comparable facility in another country. These international teams need communications, data sharing, and collaboratory (i.e., long-distance experimental collaboration) capabilities equivalent to those of a team confined to one country. A problem arises because networking is usually provided on a national basis to meet the needs of the providing country, but the same level of planning and resource provision is not directed to connecting the networks of individual countries. ESnet approaches this problem by

- Making bilateral agreements with research networks in countries most vital to DOE research, often sharing the costs of international links with those other countries.
- Organizing approximately annual meetings of a loose coalition of research networks and institutions that find they need to work together to meet the needs of international research.

In addition to the question of responsibility, 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 possible network configurations.
- Within DOE, the different research programs have connectivity requirements that vary country by country, and this makes the process of prioritizing requirements within the user community cumbersome. (Program requirements for international bandwidth are discussed in Chapter 3.)
- Traffic from one country to another may pass through a third country, e.g., from Germany to Japan via the United States, or from the United States to China via Japan.

ESnet Implementation Strategy

Overview

The responsibility for implementing the ESnet vision is shared among many elements in the DOE science community. These elements play different roles as technologies and applications migrate from the R&D stage through widespread deployment. Since program requirements often push beyond the current state of the art, an active research effort within the community is essential. This research is carried out by ESnet staff, by other groups funded by MICS, and by groups at laboratories and universities funded by other DOE science programs. Of course, the staff also tries to take the maximum advantage of developments funded by other agencies and in the industrial sphere, often working in partnership to introduce results of the latest research into commercial products.

Working closely with the ESnet staff, the Coordinating Committee (ESCC) plays a crucial role in integrating new technologies into the site infrastructures and in developing the necessary operational experience. The ESCC also works as an advocate, encouraging deployment of advanced technology and services. As services move into full production mode across the network backbone, they are deployed and supported by the central ESnet team. The most important of these services are purchased bandwidth and switchgear, which underlie the entire network. Finally, management and oversight are provided by ESnet management, DOE, and the national laboratories.

In many high-tech areas, determining a viable strategic implementation plan can be difficult because of the rapid pace of technology development. This is particularly true in the area of networking for the following reasons:

- The rate of "technology churn" is exceptionally rapid: the typical "half-life" for a technology is said to be roughly 18 months.
- The rate of performance growth to meet demand in networking exceeds that of Moore's law for computing. For example, the rate of growth for ESnet traffic has been 100% per year for more than 10 consecutive years.
- Networking in the R&D field is heavily, and probably uniquely, dependent upon the voluntary cooperation of and interoperation with many organizations (i.e., peering or interconnections). For example, ESnet has well over 100 peering agreements. It has often proved wise to judiciously accept, adopt, and/or adapt to unplanned "opportunities" for interconnection when they arise.

The ongoing implementation plan is by nature short term. The real long-term implementation plan is "be nimble, be flexible, move fast, and jump on opportunity." The ESnet implementation strategy is based upon these principles and on the elements described in the next section, which have proved to be essential to the long-term success of ESnet.

Approach

Centralized and Distributed Approach

The ESnet Project uses a combination of centralized and distributed resources. The ESnet Project, which is centrally funded and staffed, provides networking-related services to the ESnet community. Sites connected to ESnet provide the local area networking internal to the laboratory. Operation of the production network requires the cooperation of both ESnet staff and local-area network staff at the sites, and both groups discuss research and advanced deployment through the ESCC.

Extensive Collaboration

The ESnet Project is structured as a collaborative effort involving project staff, vendors, the DOE community, academia, commercial organizations, and international partners. The ESnet Project provides "core services" such as high-performance, wide area network resources with round-the-clock trouble resolution. However, the success of the effort depends on cooperation and collaboration among many people: the ESnet staff; vendors of equipment and communications services; the DOE network technical, engineering, and research communities; and members of the academic, commercial networking, and international communities. The extent of this cooperation is unprecedented and requires ongoing attention to the sociological aspects of networking as much as the technical aspects. ESnet may easily be the broadest cooperative effort in DOE.

Role of ESCC

One of the ESnet Coordinating Committee's primary functions is facilitating the introduction of new networking technologies at the sites and facilities served by ESnet. These activities are responsible for a measurable part of the effectiveness of ESnet in providing transparent, low-latency, high-performance connections.

DOE's advanced applications require access to new features and services that will require both hardware and software coordination across ESnet and into laboratory site networks. To satisfy the end-to-end needs of these applications, hardware procurement and software upgrades frequently need to be coordinated with ESnet staff and laboratory sites. The ESCC task force structure is well suited to this role.

This ESCC task force model applies equally well to delivery of advanced technology for any of DOE's other advanced applications. These include, but are not limited to, future IPv6, voice over IP, collaborative immersive visualization, and latency- or jitter-sensitive real-time applications such as computational steering. Finally, applications requiring encryption or other special services (such as group key distribution for encrypted multicast traffic) will require the coordinated distribution of controlled software through a trusted chain of custody, which a task force can also provide.

The Access Grid project is developing the next generation of remote collaboration tools, relying on high-speed networks such as ESnet

At the same time, ESCC working groups such as the Distributed Computing Working Group (DCWG) or Remote Conferencing Working Group (RCWG) will play a continuing, long-term role. For example, as DOE focuses more effort on computational grids, the DCWG will distribute libraries, serve as a local interface for security issues, and generally act as a point of contact for users. Similarly, the increased emphasis on video, particularly high-quality group video, will make the RCWG a focus for some time to come.

Role of ESRSC

As is mentioned elsewhere, the ESnet Steering Committee has recently chartered a new subcommittee, the ESnet Research Support Committee (ESRSC), that will focus on the network research needs of the DOE user community and ESnet. As research recommended by this committee and sponsored by the SC becomes ready to move from an R&D environment to the production environment, the Coordinating Committee will play an active role in facilitating this move.

The ESRSC formalizes the role of the ad hoc task force that was created in 1998 for the DOE Science Grid testbed. This formalization emphasizes the increased importance and value ESnet attaches to a persistent testbed service. This service will support experiments by network researchers and by middleware and application programmers. It will also allow ESnet and laboratory support staff to gain valuable preproduction experience. A testbed infrastructure will provide the ability to set up and use different testbed services that change as experimental needs change. By using such testbed services, the DOE research community can design and deploy long-running experiments that allow applications to use advanced network services. ESnet staff will also be able to test new wide-area networking technologies that will be used to build future production networks.

Currently, the testbed service involves four ESnet sites. It offers a quality-of-service (QoS) testbed based on a differentiated services architecture of the type recommended by the Internet Engineering Task Force (IETF). The testbed operates in the following way. User applications at the sites indicate their priority needs to host-based middleware. This middleware communicates the priority information to a reservation manager, and the reservation manager communicates with the site's routers to tell the router which packets to mark as high priority and which to leave as low priority. Routers supplied by ESnet then perform priority queue management functions to forward high-priority packets before low-priority packets are serviced. This policing function ensures that the high-priority packets receive the level of QoS the application needs. ESnet staff, site staff, and researchers work together closely on this experiment to determine what problems exist and test possible solutions.

Future testbeds will be created as necessary to experiment with grid services (i.e., security, authentication, and authorization), or to meet specific application needs.

The ESnet Research Support Committee will actively facilitate the introduction of new protocols and services into the ESnet persistent testbed infrastructure. In turn, the ESnet Site Coordinating Committee will actively facilitate moving these protocols and services from the testbed to full production.

Commercial Services

ESnet has long based its implementation approach on a combination of commercial services and "in-house" capabilities. Roughly half of the annual budget is spent on commercial services. ESnet personnel have typically maintained responsibility for the so-called layer-three services (e.g., the IP layer), while commercial services are procured to provide layers one and two. This combined implementation approach has proved advantageous.

Vendor support. ESnet has typically selected, through competitive procurement, a single long-haul communications vendor with which to work on a collaborative basis, using a multi-year master contract. This arrangement facilitates the development of a working partnership, while allowing ESnet to incorporate and use advanced communications technology as it becomes available.

Sprint has served as the selected vendor since August 1994, providing ATM services at speeds ranging from T3 (45 Mbps) to OC-12 (622 Mbps). This early solicitation by ESnet for nationwide "fast packet services" resulted in the first large-scale deployment of ATM. This step has been called the catalyst that helped crystallize the industry commitment to ATM technology. The contract with Sprint was extended and now will expire in August 2001.

The ESnet Project recently completed a competitive procurement action that resulted in the selection of Denver-based Qwest Communications International, Inc., to provide high-performance network services. The contract can extend for a maximum of seven years, beginning in December 1999. This procurement action was taken in anticipation of meeting future high-performance requirements and will establish a new architecture for the next generation of ESnet, known as ESnet 3. Initially, Qwest is to supply high-speed ATM services, but the contract allows for other advanced services (some of which may not yet exist). The contract also allows for testbed-related services and collaborative research efforts with the vendor.

Core competence. The technical staff of ESnet is a cadre of professionals who are both competent in network engineering and well versed in the networking requirements of the DOE research community. Maintaining a pool of technical talent with this rare combination of expertise ensures that ESnet can select the most appropriate combination of provided and procured services while operating production-quality services with state-of-the-art technology.

DOE-specific approach. The strategy and implementation of ESnet are guided by the needs of scientific research in DOE, rather than by the agenda of a vendor, which typically would be to optimize short-term profits.

Cost-effectiveness

ESnet has been faced with a growth rate in network traffic of 100% per year for more than 10 consecutive years, while the budget has been constrained to a much slower growth rate. This has required ongoing attention to cost-effective approaches for the project, including those noted in the following paragraphs.

The report of the 1998 program review of ESnet noted that "ESnet cost-effectiveness is outstanding, and this suggests that any increases in the ESnet budget should be viewed as among the most solid of potential DOE strategic investments."

Volume and special discounts. The ESnet Project procures services and equipment on behalf of the entire community. This volume of business can result in significant discounts. Additionally, because of the emphasis on using emerging technology in a production environment, special discounts are offered by vendors wishing to encourage the rollout of new technology.

Hubbing arrangements. Although the DOE research community comprises many sites that are "off the beaten path," there can be significant cost advantages to combining the network traffic for geographically collocated sites (e.g., in the San Francisco Bay Area) into a central hub site. The hub then carries aggregated traffic to and from a commercial service.

Local loop cost reduction. Local loop access costs are proportionally very expensive and consume a significant fraction of overall communications costs. The ESnet master contract has provided sufficient incentive for vendors to install new local loop facilities into major sites and to significantly reduce monthly local loop costs where feasible.

"Free" peering arrangements. The global collaboration requirements of the DOE scientific community require broad connectivity to the global Internet, both commercial and R&D. ESnet maintains direct interconnects (peerings) with more than 100 other networks. The cost of establishing and maintaining these peering connections can be prohibitive. However, ESnet has emphasized use of public peering locations for some time. By establishing a single connection to a public peering location, many peering connections can be established at one time.

International collaboration. ESnet represents a single entity with which international R&D collaborators are interested and motivated to establish cost-effective interconnections.

Common approach. ESnet provides a common approach to wide-area networking for the community. This common approach provides a number of advantages, including greatly increased opportunity for collaboration among ESnet sites, with associated reduction in overall effort.

Centralized networking services. The ESnet staff does planning, problem resolution, traffic management, peering agreements, procurements, installations, and so forth on behalf of the entire community. This centralization prevents duplication of these services and the associated effort at each site, thereby reducing costs substantially.

Connectivity

For external connectivity, ESnet will emphasize connections to major national peering points, including FIX-W, MAE-W, Chicago, NY NAP, MAE-E, PacBell NAP, and others if they arise. However, private peering may become necessary in the near future.

For access to most major academic sites, ESnet will continue to rely largely on Internet2/Abilene. Existing peering connections with Abilene will be upgraded. In addition, direct ESnet connections to university "GigaPoPs" will be established or maintained as appropriate and consistent with mission requirements. Access to academic institutions not covered by these two approaches will be handled on a case-by-case basis.

Development of international connectivity will focus on the "big seven," that is, Canada, CERN, England, France, Germany, Italy, and Japan. The target plan calls for ESnet to achieve a capability of approximately 5 to 10 Mbps to each location over the next few years. In addition, some programs are specifying substantially higher international requirements. The ESSC will review these specifications for priority and possible joint funding.

Budget

The importance of ESnet to DOE science is reflected in its long history of growth and corresponding budget increases of about $1.5 million per year based on the last 10 years.

FY 2001 budget breakout

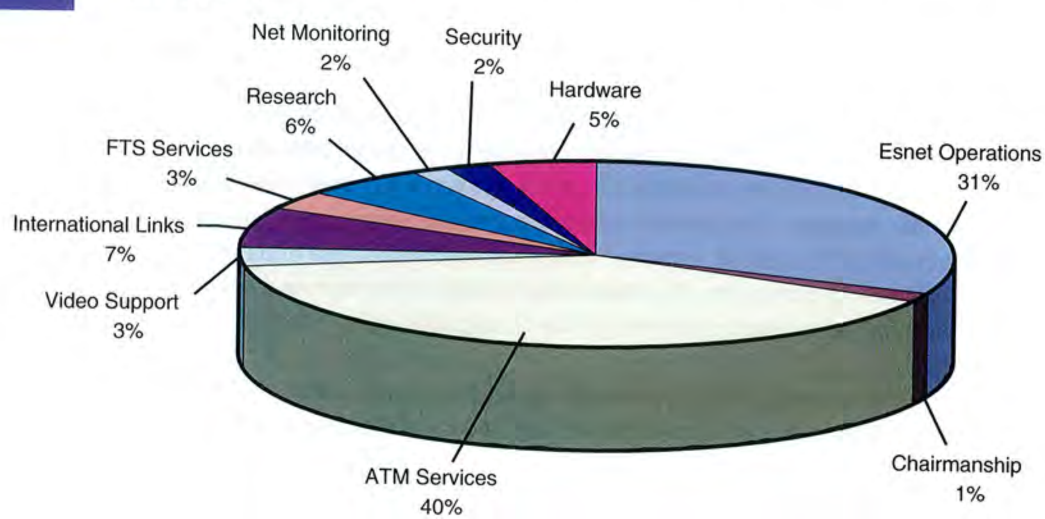

As the need for better international collaboration has grown, DOE has increased the funding for work to facilitate connectivity on an international scale (see chart). The international links available through ESnet are an important resource for the DOE research community. In addition to providing effective and cost-efficient support for collaboration, these links have served as a research tool to explore the difficulties of connecting ESnet to the various international domains and networks. Close cooperation and communication among all the communities involved will continue to be central to the success of ESnet and the mission of DOE science.

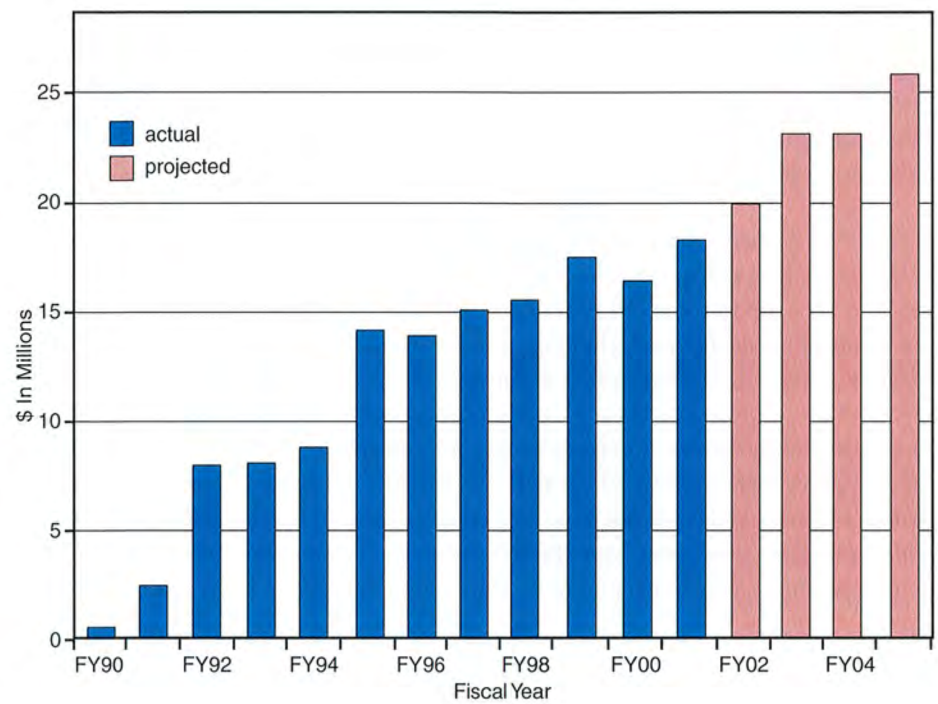

ESnet total budget

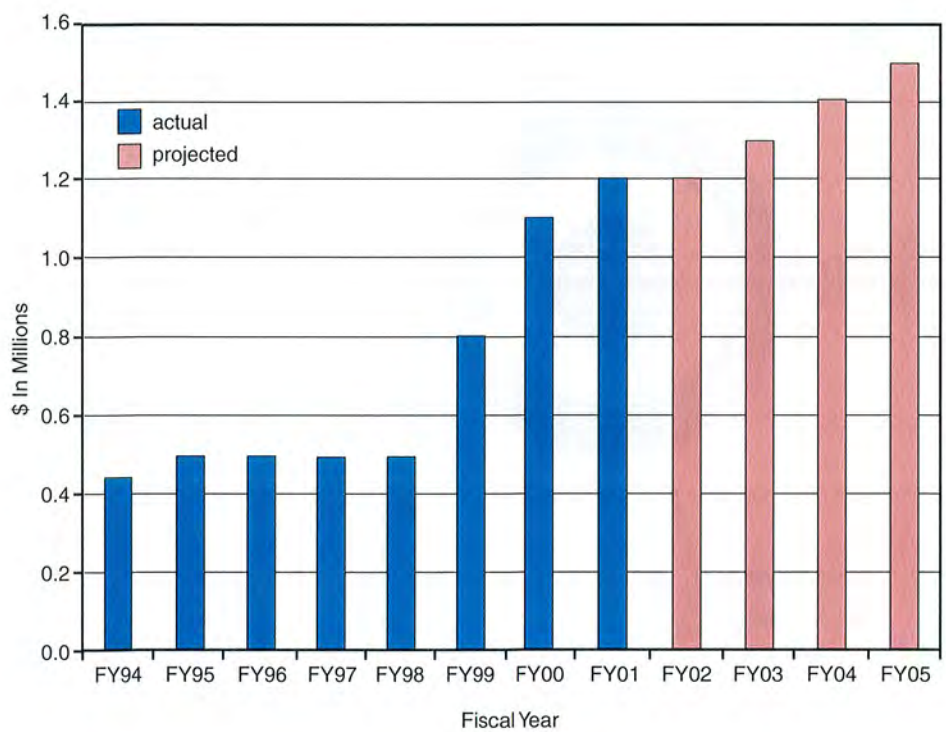

ESnet international budget (included in the total budget shown above)

ESnet 3 Architecture

The network backbone architecture for the next generation of ESnet (ESnet 3) is shown in the figure below. This architecture was chosen to minimize ongoing operational costs while maintaining the flexibility to add new technology in the future without major disruption to the operation of the network.

The design is based around five major "hubs" located in geographic areas where access for multiple ESnet sites can be consolidated. The connectivity is provided by Qwest under a new contract. Each hub is a major access port into Qwest's long-haul facilities, so this expensive component can be sized for, and shared among, several sites. This approach substantially lowers overall costs. Three smaller hubs are also implemented using the same concept. However, these "sub-hubs" then connect through the nearest major hub, rather than directly to the Qwest facilities. With this hub-based architecture, the local access to the hub for each site can be selected on a case-by-case, cost-effective basis, independent of the long-haul technology.

Flexibility is achieved because the technology connecting the hubs is independent of the technology connecting each site. For example, in the current implementation, the long-haul communications technology is ATM, while sites are connected to the nearest hub via various technologies, including SONET (synchronous optical network), ATM, T3, T1, frame relay, and private fiber.

The initial choice for inter-hub communications was ATM at speeds of OC-12 (622 Mbps) or higher. ATM was chosen (1) to ease the transition from the incumbent vendor to Qwest and (2) as the most cost-effective means of meeting the nation-wide high-speed demands of the ESnet community. The architecture allows this component to be changed in part or in total as program requirements rise and

ESnet 3 initial configuration, top-level view

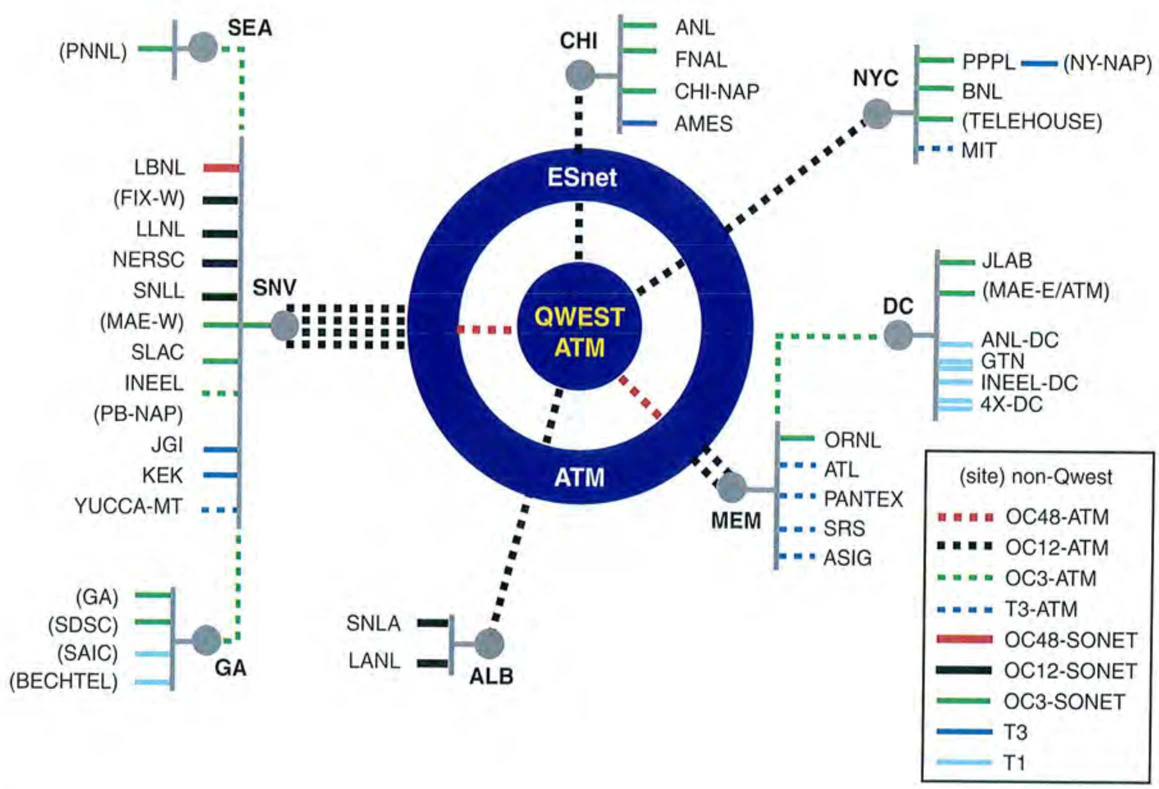

the cost per bit of other technologies drops (e.g., for packet over SONET [POS] or dense wave division multiplexed [DWDM] transmission), changing the relative cost-effectiveness of ATM.

Finally, the architecture includes a layer of "ESnet ATM" on top of the public ATM supplied by Qwest. This layer allows the various access connections to be consolidated into the single Qwest port. However, the private ATM layer also allows local communications between sites within a hub to be managed without using the Qwest port. In other words, a nontrivial portion of requirements can be met without increasing the long-haul communications requirements and the associated costs.

Research

ESnet has strongly emphasized the importance of the "production quality" of its support, balanced against the need to emphasize leading-edge technology. This approach also creates the need to integrate applied network research and advanced technology. Current and projected requirements call for network research in the following areas.

Quality of service (QoS). QoS technology provides a means of implementing priority service for specified traffic on an end-to-end basis. For example, videoconferencing traffic could be sent with higher priority than e-mail traffic. Although the concept is fairly simple and has been tested successfully in the ESnet testbed environment, a substantial number of other issues must be addressed to fully support QoS capability. These other issues include traffic policing and shaping requirements, allocation, accounting, authorization for users, deployment in each site's local-area network, and diagnostic tools. Limited initial deployment is planned for FY 2001.

High-performance technology. The ESnet Project has an ongoing interest in cost-effective, high-reliability, high-performance architectures. At the same time it maintains an interest in state-of-the-art communications services. Current leading-edge technologies that define such services include IP/SONET (direct IP over SONET) and DWDM.

Multimedia support. Multimedia support will become even more important to research collaboration. IP-based videoconferencing is maturing quickly, and research into such H.323-based capabilities is in progress. Additional areas of interest include voice over IP and video streaming.

Security. Research areas that could address security concerns for the network include virtual private networks (VPNs) and policy-based routing.

Advanced protocols. A new approach called multi-protocol label switching (MPLS) may be of significant interest to the DOE research community. MPLS offers high-performance switching, source-specified routing (e.g., QoS and testbed routing), and traffic engineering.

IPv6. ESnet will remain in a leadership position in testing and deployment planning of IPv6, which is already supported by ESnet.

Additional Services

The ESnet Project continually considers additional services that could be provided by the project. To date, the major addition to the backbone network services has been digital collaboration services (DCS), a collection of collaboration capabilities that may be reserved through a central, Web-based reservation system. New capabilities, including H.323 (IP-based videoconferencing) and streaming video, are being evaluated as candidate components of DCS.

The emergence of computational "grids" creates a need for various support services that may be appropriate for ESnet. Currently, project staff are evaluating directory and naming services. They are also studying the issues involved in establishing production-level capabilities for certificate authority and certificate servers based on public key infrastructure (PKI).

Cyber Security

ESnet is a federally funded resource and therefore has requirements to ensure that it is used appropriately in compliance with its funding. A document outlining the ESnet Acceptable Use Policy is available online *(http://www.es.net/hypertext/esnet-aup.html).*

Although the physical resources of ESnet are distributed throughout the country, many of these resources are located at LBNL, along with the project office. For these reasons, the security policy and practices of ESnet are in concert with those of LBNL (see *http://www.lbl.gov/security/*) as well as the applicable requirements of DOE.

ESnet is chartered to support exclusively open research and transport unclassified data, and therefore the security issues are those associated with an unclassified facility. Management responsibility and control extend to the "edges" of the network. The connected sites and their resources and computers are not under the purview of ESnet, and site security is considered a local responsibility. Within ESnet, responsibility for cyber security covers the unclassified security issues associated with protecting ESnet resources from inappropriate or malicious use.

The ESnet Project does not assume or accept responsibility for security at connected sites or the privacy or integrity of data in transit across the network. However, the ESnet architecture provides an initial level of protection for sites. In addition, activities are underway to support end-user requirements for a higher level of security than is available over an open, unclassified network.

The unclassified research programs within SC need network and collaboration support with both state-of-the-art performance levels and global interconnectivity. This double requirement increases the difficulty of implementing strong security, even in an unclassified environment such as ESnet. These three divergent issues—security implementation, performance, and connectivity—must be kept in balance while working to meet user needs.

Security technologies continually evolve to keep pace with the rapid changes in both networking technology and intrusion methods, and ESnet staff will continue to evaluate new security technologies as they become available.

2 ESnet Futures

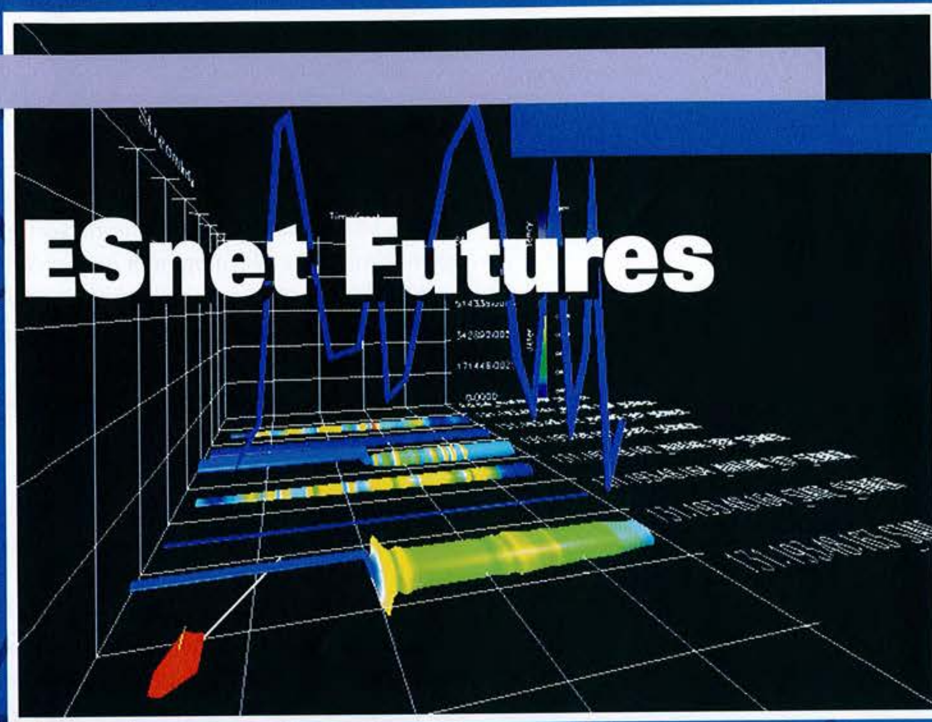

Vision

Advanced network communications are increasingly critical for accomplishing DOE's science missions. Development of new computer communications technologies, both within the DOE community and elsewhere, are creating outstanding opportunities for scientific advances and technical leadership. With these opportunities come many challenges in which rapidly growing demands for network capabilities play a major role. How those demands are addressed will shape DOE research in very significant ways.

Historically, DOE has been a leader in developing and using large-scale computing systems and applications for science. Today's most demanding problems continue to stretch the limits of large-scale computing, data storage, and networks, and there are strong indications that growth will be rapid in the next five years. Among established research programs, a wide range of science and engineering pursuits are preparing to move "off the desktop" as researchers strive for greater accuracy, definition, and complexity in their models and simulations. This change also reflects the greater importance that modeling and simulation now have in science. In general, emerging trends that will affect ESnet include informatics, entirely new uses of computing, integrated environments for scientific collaboration, distributed resources, and dynamic connectivity. These trends appear to be converging in what is sometimes called the DOE Science Grid.

Informatics. Sophisticated informatics (i.e., data acquisition, management, and analysis) is now broadly viewed as an essential capability in cutting-edge instrumentation. New classes of instrumentation are emerging that require real-time supercomputing to analyze data while the experiment is going on, so that decisions can be made about the direction of that experiment. As examples, imaging of living cells under stress and determining protein structure from high-resolution mass spectroscopy both require rapid decisions about transient processes and species.

New uses of computing. New scientific challenges are also driving the development of wholly new computational thrusts. The informatics revolution associated with genome research is an example of large-scale requirements emerging where none existed before. Altogether, a much larger and more diverse body of high-end supercomputing and network users is forming. These new users will not only need terabyte-scale computers and tera/petabyte-scale archives, they will also be installing an entire spectrum of midrange systems and databases. Overall, the expansion of the supercomputing community can be expected to drive ESnet requirements up at rates exceeding the Moore's Law growth of microprocessor performance.

Distance collaboration. Advances in computing and communications technologies are also enabling scientists to work together at a distance. DOE collaboratory programs are providing new technologies and proofs-of-principle that are changing the design of many next-generation science programs. Today, a small number of domain-centered collaboratories enable research at a distance and interaction among colleagues. These collaboratories will proliferate, especially within the scientific communities around the unique resources centered at DOE's user facilities.

In the next stage, cross-disciplinary collaboratories will arise to address complex interdisciplinary problems related to DOE's missions, for example, collective phenomena, environmental remediation, "green" energy systems, cellular response, nanotechnology, and carbon management, to name just a few. Experience with these collaboratories will lead to novel uses not envisioned previously. Collaboration capabilities will also fuse with application user interfaces and problem-solving environments to create integrated science environments with built-in collaboration capabilities. Hence, during the next five years, collaboration technologies will move from the cutting edge to a core network and application requirement for competitive

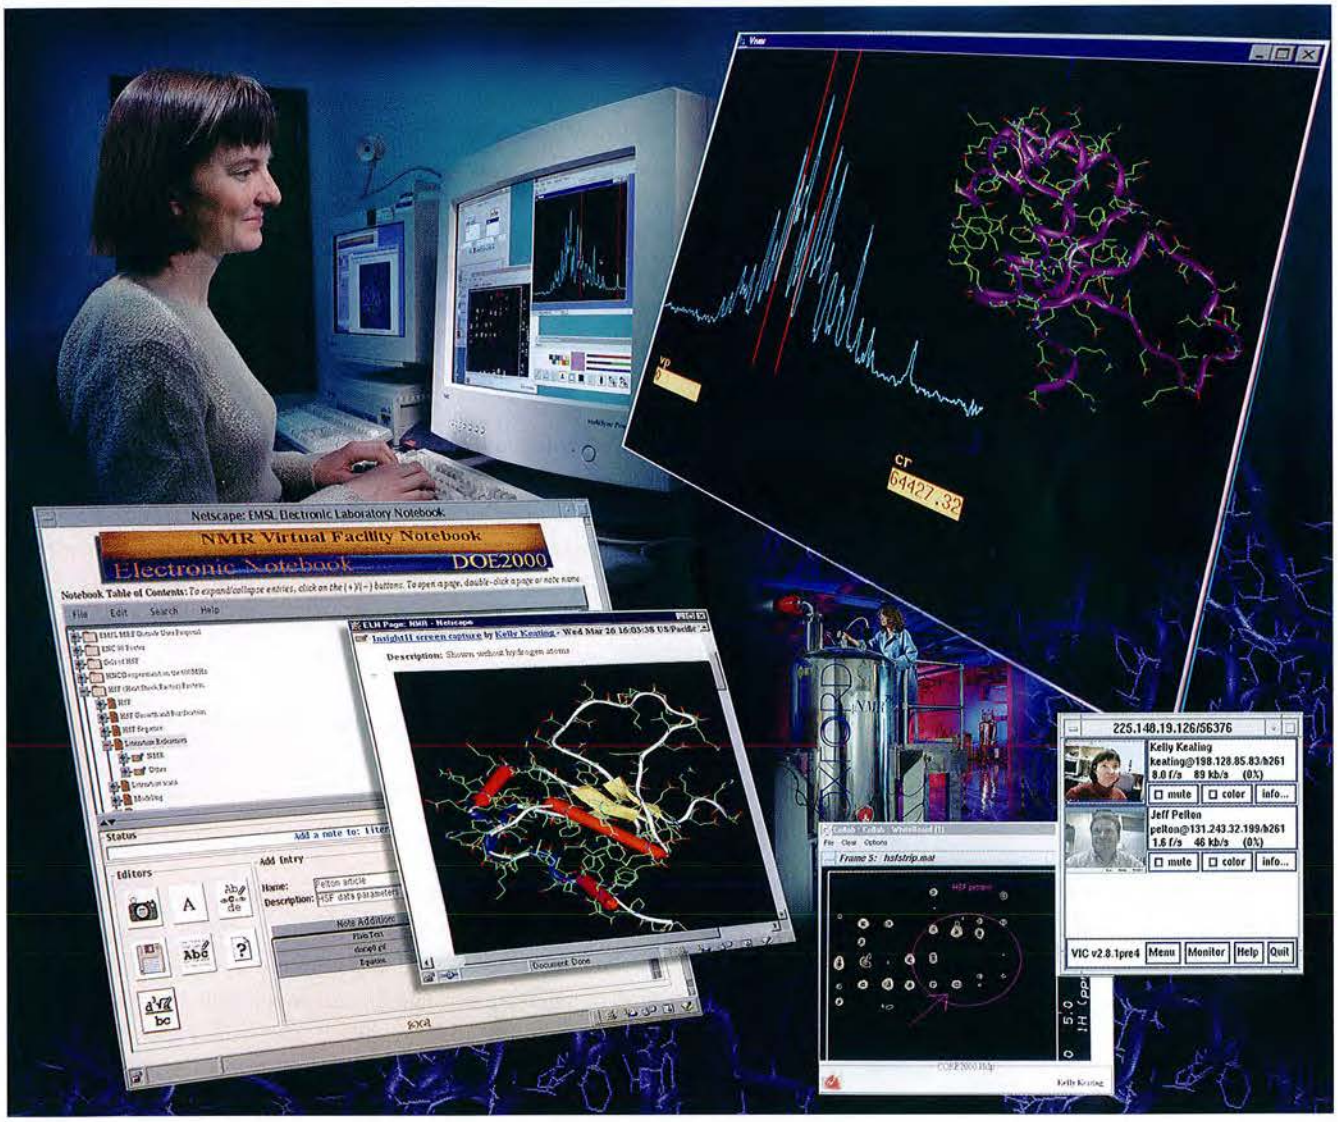

Virtual Nuclear Magnetic Resonance Facility at PNNL's Environmental Molecular Sciences Laboratory (http://www.emsl.pnl.gov:2080/docs/collab/EMSLVNMRF.html). The facility allows remote researchers to run spectrometers, consult with staff, and share notes and data, all via the Internet. It is based on tools developed for DOE's national collaboratories program.

research. Such collaboratories will feature a diversity of data streams, each with its own delivery, reliability, and security requirements to be supported by ESnet services.

Distributed resources. Today, an individual researcher uses only a relatively small number of laptop, desktop, server, supercomputer, database, and archive resources that are largely directed by the scientist from a single point, the desktop. However, large investments in middleware, coupled with advances in network performance and services, will soon enable seamless distribution and federation of scientific resources, increasing the number of systems routinely employed by scientists.

Dynamic connectivity. In addition, computing and information access points are poised to move off the desktop; that is, place-based computing will be replaced in many instances by location-independent computing. As wireless local-area network and wide-area network capabilities become ubiquitous, a plethora of specialized mobile devices will complement the "one box does all" personal computer. Individuals will compute and collaborate wherever they are. Wireless devices may proliferate faster than wired devices, adding not only devices to networks, but also new classes of dynamic connections with particular ESnet service requirements.

DOE Science Grid. Many of the trends described above serve to knit together the collection of researchers and resources used for DOE science. DOE is making a transition from a collection of laboratories to a seamlessly integrated scientific computing enterprise. In this enterprise, sometimes called the DOE Science Grid, all major computing and data resources will be securely available from any DOE researcher's desktop.

A system of interoperable credentials will provide access to resources across the enterprise and beyond in the national and international science community. If planned with foresight and cooperation, such allocation, authorization, and security processes will interoperate across agencies and industries. Databases will transition from simple repositories to qualified information resources. Many will be "smart" about how they use the network, and some will be able to dynamically tap distributed resources to satisfy requests.

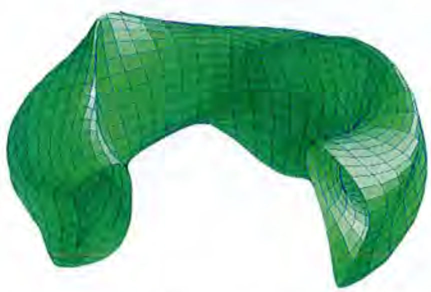

Computer model of magnetic surfaces in the proposed National Compact Stellarator Experiment (NCSX) (http://www.pppl.gov/ncsx/). These computations are the product of a team based at national labs and universities.

In this complex environment, the health of the network is paramount. Proactive network monitoring and predictive intrusion detection will allow network engineers to make more informed decisions and avoid problems that would impede work or threaten security.

Implications for ESnet. Hence, in the next five years, ESnet will support a more diverse portfolio of capabilities for a more diverse portfolio of science. Robust and reliable production services are essential, as well as a commitment to network research for next-generation capabilities. Bandwidth usage will rise considerably faster than Moore's Law, driven by a host of extremely challenging research problems, the broader penetration of terascale computing, tera/petascale databases, and the integration of collaboration and remote access technologies into broad practice. Seamless integration of capabilities and services from many sources will be needed, incorporating both community and commodity components and application interfaces. Such network capabilities are at the heart of enabling the future vision of ESnet, offering an entire spectrum of capabilities, from palmtop to grid computing. The rest of this chapter reviews the directions in which ESnet will need to move to support these trends while maintaining its high level of service: advanced applications (collaboratories and grids), infrastructure development, and research.

Advanced Applications

This section discusses application areas that, with appropriate development, promise to serve as a foundation for expected expansion and extension of DOE research programs. With resources and researchers scattered around the country (and the world), collaborations of scientists are typical in the DOE science community. In the high energy and nuclear physics areas, as well as in the fusion community, for example, collaborations can encompass hundreds or thousands of scientists who rely heavily on effective interactive, network-based communications.

These groups and many more smaller collaborations need to share access to scientific instruments, data, and computational resources. Collaborative authoring of research papers, computer codes, and other written documents by extended groups is an early and continuing use of the network. In addition, remote conferencing is becoming more important to all aspects of the research effort. In the following discussion, the major applications of advanced network technologies are described. These are divided into four rough categories:

- Collaboratories (distributed collaborations).
- Computational grids (distributed computing).
- Data grids (distributed databases).
- Visualization grids (distributed visualization).

Collaboratories

Many of the facilities of interest to the ESnet community are unique and too expensive to duplicate. Traditionally, this meant that experimentalists were required to travel to carry out their work. In recent years, a new approach has emerged, one that relies on high-speed networks to allow remote access and remote control of experiments. This approach reduces travel time and brings the maximum manpower to bear on scientific problems. At the same time, computational scientists, who have long relied on the networks for access to supercomputers, have begun working in geographically distributed but highly coordinated teams. Both groups require an infrastructure that provides network connectivity and a range of services to support the use of widely distributed resources. At a minimum, the infrastructure must support transparent access to data, shared applications, audiovisual communication for operations and planning, and security.

Vigorous development of collaborative environments will bring substantial benefits to research programs, including more efficient utilization of costly or unique instruments, as well as researchers' time. Further, the ability to correlate results from multiple instruments will greatly improve the sophistication of scientific insights resulting from such measurements. The following discussion outlines several issues for future development of collaboratories.

Experience to date. While remote operation of scientific instruments and facilities remains an experimental capability, the ESnet community is rapidly gaining experience in this area. Fusion energy researchers have demonstrated the practicality of this approach by running major facilities from remote sites, for example, by running a tokamak at the Massachusetts Institute of Technology from a control room at Lawrence Livermore National Laboratory in California *(http://www.psfc.mit.edu/people/g/papers/remote.html).*

A number of well-organized efforts were spawned by the DOE2000 program, which was an effort to develop new computational tools to advance the concept of "national collaboratories." These projects included the Diesel Combustion Collaboratory *(http://www-collab.ca.sandia.gov/dcc/research.html),* the Environmental Molecular Sciences Collaboratory *(http://collaboratory.pnl.gov/),* and the Materials Microcharacterization Collaboratory *(http://tpm.amc.anl.gov/mmc/),* for example.

Needed development in support of remote experimental operations requires close cooperation between network managers and software developers on the one hand, and instrument builders and users on the other. Experience has shown that remote participation is most efficiently implemented if it is designed in from the beginning.

Nonlinear simulation of magneto-hydrodynamic instabilities driven by strong gradients in plasma pressure

Data transfer. A key element in running an experimental collaboratory is the ability to transmit and display large amounts of data from the experiment to the remote team (or teams). Remote researchers must also have access to shared applications—those that support viewing and interaction from two or more sites simultaneously. Important components in this environment include access to large relational and object-oriented databases; the capability to create, access, and use electronic note-books; and visualization software and tools for analyzing application performance. All of these requirements need to be met in real time.

Simulated environments. Experience has shown that an "immersive environment" similar to the environment at the experimental site, including excellent audio and video communications, can be a crucial element in making remote participation a success. The goal is to allow users to see and hear each other, to virtually "touch" and operate a representation of the instruments, and to visualize data—all in a simulation of the actual environment.

This simulation may be crude, using low-performance streaming audio and video in addition to data display tools, or it may extend to virtual reality display environ-ments at each location operating through high-performance network connections. The resulting data flows will vary according to the degree of simulated reality that is required. Diagnostics, remote access to machine status, and other related sub-systems must be supported by the network with appropriate prioritization.

Security. As researchers gain the ability to remotely control expensive facilities in part or whole, the need for strong authentication and security services becomes essential. These services are required at some level for all networked applications but are particularly important for experimental collaborations where unique, costly, and possibly hazardous equipment may be involved. The remote control of experi-ments 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.

Code management. Experimentalists are not the only customers for these new capabilities. As the power and complexity of computational science grows, codes are increasingly the products of large, geographically dispersed teams. While this has brought new resources to the development and use of codes for simulation, data analysis, and other purposes, it brings with it additional requirements for the computing framework and for coordination and planning.

The framework must support collaborative and distributed code development, distributed computing and visualization, documentation, shared access to results, and shared logging of runs. The following are among the issues that need to be addressed for progress to occur: efficient use of distributed compute cycles, reliable asynchronous intertask communications, multicasting of data, remote display and downloading of results, distributed task queuing, and session management. Planning and coordination require the deployment of various modes of videoconferencing, a common version-control system for code, an electronic notebook to document coding changes, and other available communication and collaborative technologies. A common online code-sharing library, including essential middleware, would improve code and data access, code interconnection, and code invocation.

Computational Grids

With the widespread use of local-area networks, computing moved from a centralized to a distributed model. It is common for system files, user files, data archives, relational databases, and applications to be spread among different machines and shared by users who can use them all as local resources. In more recent years, this model has moved onto the wide-area network and has grown to include supercomputers, computational clusters, massive high-performance storage systems, and other specialized resources. The DOE Office of Science computing community forecasts that in the next five years, the flagship computer for DOE computing will have a capability of as much as 100 teraflops (i.e., trillion floating-point operations per second). In addition, powerful local workstation clusters have been developed to provide supercomputer-like speeds.

The goal of a computational grid is to create an environment in which users share and configure these powerful resources into a single virtual computing engine and thus reduce the amount of "wall clock" time required for calculations by orders of magnitude. In this sense, the grid is an enabling technology: without it, many calculations would not be done. Large consortia funded by DOE and other agencies are working on this challenge. Two such organizations are the Global Grid Forum *(http://www.gridforum.org/)* and the Globus Project *(http://www.globus.org/)*.

At one extreme, such an environment would allow the sharing of a few distributed and unique high-performance computers. At the other extreme, the environment might consist of many underused desktop machines that could be shared. Finally, between these two extremes, a new class of distributed computing is growing, based on coupling the resources on a user's desk with a range of computing resources that are physically located both at the user's site and at remote sites.

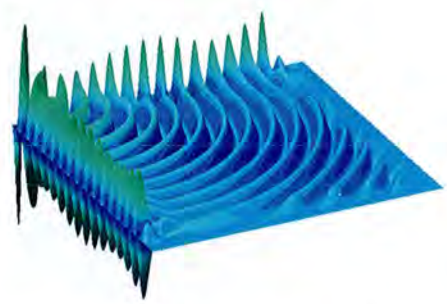

Wave functions for the breakup of a system of three charged particles illustrate the first complete solution to a fundamental problem in quantum physics, calculated at the National Energy Research Scientific Computing Center

As a result of this new approach, an expanded supercomputing community is forming. The associated networking requirements are also expanding as more sites assemble dedicated clusters of computers and are willing to broker time among them. Comprehensive and precise monitoring is essential for a widely distributed system that must analyze algorithms, diagnose problems, and dynamically adapt the system configuration. Active autonomous resource management will be needed for reliability. Basic management functions such as automated server restart, data migration, and congestion avoidance can be accomplished with the same type of infrastructure as monitoring. All resources must support "contract" scheduling so that all components are simultaneously available. This is essentially the problem of policy-based resource management as found in quality-of-service architectures.

Furthermore, the combination of resources applied must be matched appropriately to the problem being solved. For example, in some cases, network latency poses synchronization problems that reduce the utility of the computational grid approach. Algorithms that can successfully hide latency at the required level of tens of milliseconds (instead of microseconds) are still in the developmental stage. However, the ability to distribute problems across a meta-computer whose resources are "free" within the brokering scheme offers great rewards.

Data Grids

Just as researchers use computational resources spread across the wide-area network, they also need to access data stored in various locations. Shared file systems have existed for many years. What is new is systems that can allow transparent, shared access and present the user with coherent, complete, and well-documented data sets even if the data are physically stored in many repositories located around the world. Some of these data sets are already truly enormous, and they are expected to grow by petabytes per year. Effective use of the data puts a premium on the development of powerful and flexible visualization systems that can present selected portions of the data with the highest resolution and minimum latency.

Experiments and simulations in DOE research programs often produce huge data sets. Examples include work at the Relativistic Heavy Ion Collider, Thomas Jefferson National Accelerator Facility, and Large Hadron Collider and simulations for the Atmospheric Radiation Measurement Program and the Plasma Science Advanced Computing Initiative. For users to fully exploit their data, it will be necessary to develop architectures for high-performance, data-intensive, distributed systems. The requirements include distributed data analysis, mechanisms for managing high-speed data flows, and flexible remote access. This use of the network will require construction of an infrastructure for widely distributed analysis of data at multi-petabyte scales by hundreds or thousands of scientists. Concurrently, it will require the development of network and middleware infrastructures aimed broadly at data-intensive collaborative science.

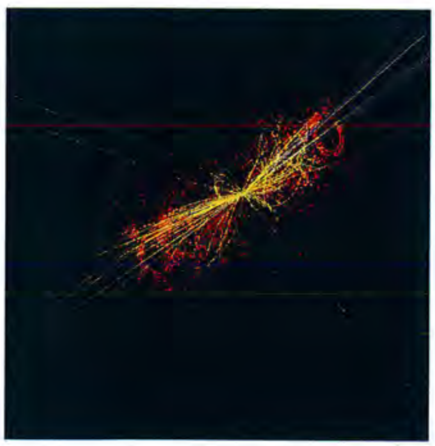

Visualization of high-energy particle interactions provided by GIOD, the Globally Interconnected Object Database (http://pcbunn.cacr.caltech.edu/)

A variety of underlying technologies exist for this purpose (e.g., distributed file 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. The choice of technologies will depend on the nature of the data, analysis, and visualization for each physical problem.

The architecture chosen for data storage also depends on the nature of the science being done. Some communities employ centralized, hierarchical systems, while others prefer highly distributed storage. In either case, distributed high-performance caching is essential. Tools to manage and integrate views of distributed data will need to be developed, along with tools for managing metadata, including data characteristics and data catalogues, all of which might be further integrated with electronic notebooks. For large data sets, these processes must be automated. Additional issues include how to provide a global name space, how and when to migrate files in the hierarchy, how to manage caches, and the choice of a data processing model.

Given such large databases, additional strategies will be needed to reduce the impact of moving the data over the wide-area network. Two such strategies are local data reduction and data compression. Methods for providing users with "cost" estimates 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. With generation of data clearly increasing faster than the bandwidth available on the Internet, it may be necessary to assess computational models in terms of their impact on the network, that is, to ask at what point the processing should be moved to the data instead of moving data to the processor.

Visualization Grids

The need for powerful visualization systems cuts across most of the experimental and computational sciences. The large data sets created by experimental facilities have already been discussed. DOE scientists routinely perform computer simulations, computer modeling, and the analysis and synthesis of large amounts of experimental data by converting data into pictures or animation. This conversion requires sophisticated, data-intensive visualization techniques. The requirements for these techniques and the associated large data sets permeate the ESnet user community, particularly for work in plasma physics, climate analysis, materials science, chemistry, computational fluid dynamics, combustion, DNA analysis, particle analysis, and astrophysics.

To be effective, a visualization environment must allow the user to select, transmit, and display large data sets pulled from enormous archives with minimal latency. Rendering engines may be collocated with the data or with the viewer or may reside at a third site. In all cases, powerful computers will be required. Applications of such scale can be executed only in environments with large amounts of memory and processor speed. An emerging trend for addressing this problem is the use of visualization tools based on massively parallel processors. Given their uniqueness and scarcity, these tools are geographically dispersed. To achieve interactive rates, images must be delivered to the desktop at 5 to 30 frames per second, challenging both the bandwidth and the responsiveness of the network. Similar requirements are typical of shared work spaces, immersive visualization environments (including latency-intolerant devices for virtual touch), and remote experimentation.

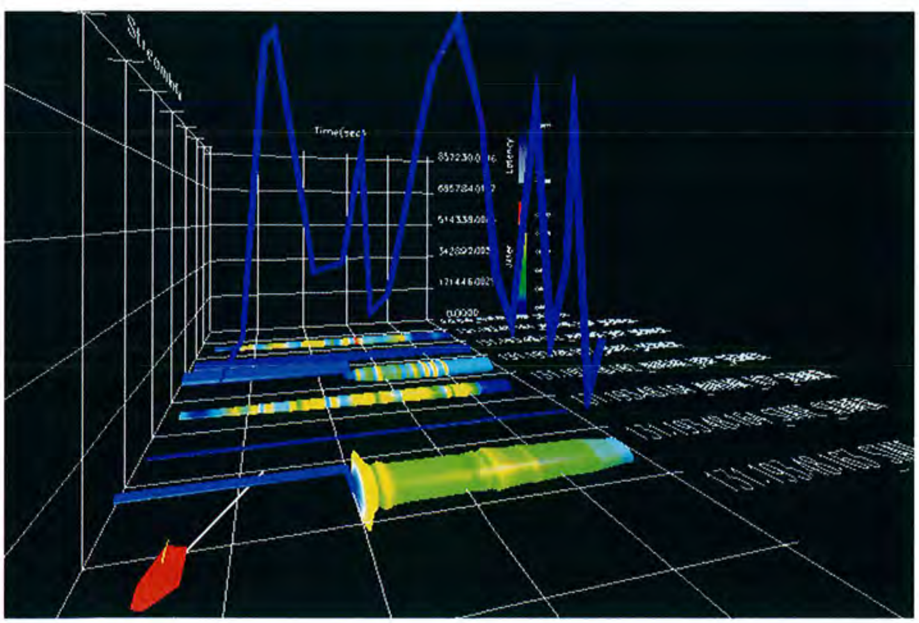

An immersive visualization of data from Environmental Molecular Science Laboratory at Pacific Northwest National Laboratory. It was built as part of an exploratory project in laboratory-scale remote operation of instruments. The data and visualization were produced by faculty and students from the University of Washington as part of a curriculum enrichment program.

Advanced Infrastructure & Services

To implement the vision outlined in the preceding section, ESnet will need to provide infrastructure and services of higher performance and far greater complexity than those available today. The common goal of these new capabilities will be to link distributed resources to researchers over the wide-area network. ESnet has been successful in providing high-speed connectivity to its backbone institutions, but to keep up with the anticipated growth in demand and to provide the foundation for advanced services, new infrastructure technologies will need to be adopted as they emerge. Improved connectivity alone will not suffice, however; the service infrastructure will take on an increasing importance. Resource management and monitoring, network transparency, distributed data access, security, and teleconferencing all require shared infrastructure. Some of that service infrastructure will be provided by ESnet directly, while other parts will be supported by laboratories and projects.

Infrastructure

The technology used in ESnet is selected carefully to provide services that maintain a balance between leading-edge technology, cost-effectiveness, reliability, and performance. When warranted, ESnet may choose to emphasize one of these at the possible expense of one or more of the others. For example, in the mid-1990s, the decision was made to incorporate ATM (asynchronous transfer mode) into the backbone because it offered possible future cost savings, although its performance and reliability had not yet been demonstrated. ATM has subsequently proved to be a cost-effective means of providing high-bandwidth communications and significantly improving the reliability of the underlying communications facilities.

The ESnet community is examining several new infrastructure technologies. Examples include IP over SONET (Internet protocol over synchronous optical networking) and DWDM (dense wave division multiplexing) at OC-48 (2.5 Gbps) and OC-192 (10 Gbps). IP over SONET allows network data to be sent directly over the SONET interconnects, eliminating overhead and complexity. DWDM allows multiple streams of data to be sent over different light frequencies (i.e., colors) on the same optical fiber, thereby increasing the overall capacity of the fiber. Optical networking maintains the data in optical form, without requiring conversion to electrical signals for processing at intermediate steps. These options represent "core" or backbone technologies: that is, they do not introduce new capabilities to the network users but instead offer new price and performance advantages that result in overall cost savings.

Other technologies being considered include MPLS (multi-protocol label switching), QoS (quality of service), IPv6 (Internet protocol version 6), and VPN (virtual private networks). MPLS attaches a short label to the network data, and this label is used to select the network path over which the data flows. This approach allows the sender to select the delivery path. QoS allows data to be marked with an indicator of its relative priority, so that higher priority data will be handled preferentially. IPv6 is an enhanced version of the protocol currently used within the global Internet (IPv4) and includes capabilities such as a much larger address space and enhanced security. ESnet is already providing an initial IPv6 testbed for testing and evaluating transition issues and strategies. VPN capabilities allow a private network to operate securely over public network facilities as though it were using its own dedicated facilities. All of these technologies can extend into the local-area network environment and will require support and coordination from end-site network managers.

Quality of Service

The success of computer networking to date has encouraged the creation of increasingly demanding network-based applications, with associated expectations about network performance. New applications will require stringent limits on latency, jitter, packet loss, and throughput. Real-time applications are expected to permeate the network as collaboratory and remote experimentation become more common. Normal network activities will also continue to grow as computers, network interfaces, and campus networks are upgraded. Given a finite budget for network connectivity, it may not be possible to "overprovision" the network sufficiently to meet all of these demands without some form of resource management. Currently, no network management mechanism is in use to allocate available resources in a manner consistent with programmatic priorities. Quality of service (QoS) is a term for a wide-ranging set of technologies aimed at matching network resources to programmatic or institutional policies.

The challenges for any QoS system are immense. Users are interested in end-to-end performance, but networked applications straddle many geographical and institutional boundaries. Reconciling these two realities will require a community-wide bandwidth broker that is cognizant of the full range of available resources. In addition, a QoS cost and enforcement mechanism must be in place across the entire system. Different applications will require different QoS implementations. For example, an application may need a specific instantaneous data rate, packet loss, jitter, or latency to be successful and would need to be able to reserve those parameters during the session. In some cases, QoS may be a static condition applying to all sessions involving remote operation of a facility. In other cases, QoS might need to be integrated with a scheduling mechanism, so that appropriate network performance can be scheduled to coincide with other reserved resources. The process of booking and then delivering a specified QoS resource is unsolved at present and is a major need of the ESnet program.

Security Services

The requirement for enhanced security for networked computer systems is all too clear today. Destructive attacks are commonplace and persistent; no institution is immune from their effects. The vision espoused in this plan—of new network-based paradigms such as distributed computing, virtual research groups, and telescience—depends implicitly on stringent security. This security includes both the end-to-end security of applications and, of course, security of the underlying network itself. The entire infrastructure must be protected: it will be an important component of the scientific environment and will be subject to the same attempted denial of service, compromise of integrity and privacy of information, and general harassment as our general computing environment.

Technologies and tools exist to strengthen the security of computing and networking environments. However, implementation can be difficult and expensive, standards or policies are often incompatible, and coordination between institutions and communities is weak. The DOE science environment is entirely heterogeneous, with national laboratories, universities, and foreign institutions working together on common problems. Security implementations need to consider this environment and protect critical resources while remaining as flexible as possible. Shared infrastructure could effectively allow the kind of cost-sharing that would bring smaller research units (those without dedicated computer security staff) into compliance.

The primary requirement in the DOE science community is for secure authentication and authorization. Resources should be transparently accessible, but only to authorized individuals. The process of authenticating these individuals should be simple, but certain. A worthy goal is the "single secure login," that is, a single login per session that provides access to all authorized resources. Needless to say, passwords or other confidential information should never traverse any piece of the network in the clear. Since research in the DOE science community is open and "public," encryption of data is usually not a goal in itself. On-demand encryption can, however, be a powerful tool for managing information that would be useful to those who aim to attack or misuse a system.

Other Middleware

The QoS and security services described above rely on "middleware," that is, software that allows an application to interoperate with other software while eliminating the need for users or applications to understand and program the low-level operations of the various parts. The vision for a DOE Science Grid depends on a wide variety of middleware services. Many of these services will be provided by commercial software, because such products have proliferated in support of electronic commerce and electronic business. Universal middleware for scientific research is unrealistic. Many middleware systems will exist within ESnet, but a critical issue will be identifying the underlying infrastructure services that are essential for all of them. Another crucial issue will be decisions on where and how to implement standards. There will not be a single model for defining, structuring, and implementing middleware: different scientific communities will emphasize different elements, because their data, analysis, and community interactions are different. However, a set of infrastructure services is needed that can support the construction of various middleware systems. The role of ESnet will be to support some subset of commercial and custom-built systems providing critical infrastructure for applications. Typical middleware services include those that support

- Client/server connectivity.
- Network transparency, including multicasting.
- Platform transparency.
- Location transparency (global directory and naming services).
- Application and tool support.
- Database management support.
- Security services.
- Messaging.
- Queuing.
- Monitoring.
- Data translation and transformation.
- Dynamic resource configuration and management, load balancing, accounting.

Remote Conferencing

Teleconferencing and videoconferencing are widely used and in constant demand in the ESnet community. Almost 300 separate conference rooms are registered for digital collaboration services. These technologies have been embraced uniformly, and rapid growth is expected to continue. As useful as they are, the tools available are still relatively primitive, and major improvements would greatly enhance their overall effectiveness. Advances are needed initially for ease of use, incorporation of multimedia interactions, shared creation and editing of documents, and integration with data collection and analysis environments. The expanded use of remote conferencing has driven the need for service directories, and the multiplicity of participants in any single session demands greater capabilities for controlling the sequence of speakers. Additionally, as tools become easier to use and are more integrated into the networked environment, improved planning and coordination services will be needed.

ESnet collaborations can be expected to make heavy use of workstation-based videoconferencing. The existing service model, however, is a barrier to widespread use, as it adequately supports only small numbers of participants and audio and video of only modest quality. Still, impressive applications can be built on this technology. The Telepresence Microscopy Site *(http://tpm.amc.anl.gov/)* enables collaborators to view operators, instruments, and data. To best serve the ESnet community, the expanded conferencing ability will need to be integrated closely with collaboratory operations. Finally, the community will need a complete and readily accessible directory of institutions and individuals who can be reached via this medium.

Network Monitoring, Diagnostics, and Traffic Analysis

Internet end-to-end performance monitoring (IEPM) is widely deployed in the ESnet community. In all, ESnet receives reports from more than 3,300 monitor-monitored pairs involving more than 70 countries, reflecting what is arguably the most extensive end-to-end network performance monitoring project in the world today. Eight ESnet sites actively run the set of monitoring tools known as PingER (for Ping end-to-end reporting). These tools were developed by staff of ESnet, the Stanford Linear Accelerator Center (SLAC), and the High Energy Physics Network Resources Center *(http://www-iepm.slac.stanford.edu/)*. Most ESnet collaborator sites are either monitored directly or are represented by so-called "beacon sites" with similar connectivity. Another 22 monitor sites are located at sites of considerable interest to ESnet (e.g., CERN, DESY, and RIKEN).

These reports provide valuable information for planning, setting expectations, and troubleshooting. They help ESnet staff set service level agreements and validate that these agreements are met. Furthermore, the data help guide improvement of the network infrastructure by highlighting common bottlenecks and important links needing upgrades and by providing a basis for long-term planning.

To keep pace with the increasing importance and extent of networking, future monitoring efforts will need to

- Extend the deployment of monitoring (in both time and space).
- Provide a scalable, manageable measurement infrastructure.
- Address increased diversity in capabilities.
- Meet new challenges imposed by security concerns.

- Provide better visualization of results.
- Provide automated, robust detection and reporting of exceptions.
- Compare measurements obtained by different mechanisms and subsequently validate and interpret the results as they relate to applications of interest to ESnet.

Furthermore, to support QoS, monitoring of pilot implementations will provide information on how well different strategies work in the face of varying competing loads for various applications. It is hoped that such studies will clarify the circumstances in which QoS is effective (e.g., for what ranges of utilization and for what applications). Of particular interest to ESnet is the competition between general traffic (e.g., Web and e-mail) and specific applications such as real-time experiment control or Internet telephony and bulk data transfer.

To support the requirements of data grids, monitoring tools are being installed at the major grid sites. These tools are also being extended to provide more accurate information needed for high-performance applications. It is hoped that the results will help researchers to make more realistic choices about bulk data transfer (e.g., network versus tape); assist the research communities in selecting sites for major data repositories, mirror sites, or regional computation centers; and provide a basis for deciding how to instrument applications and configure TCP/IP stacks to optimally use the network.

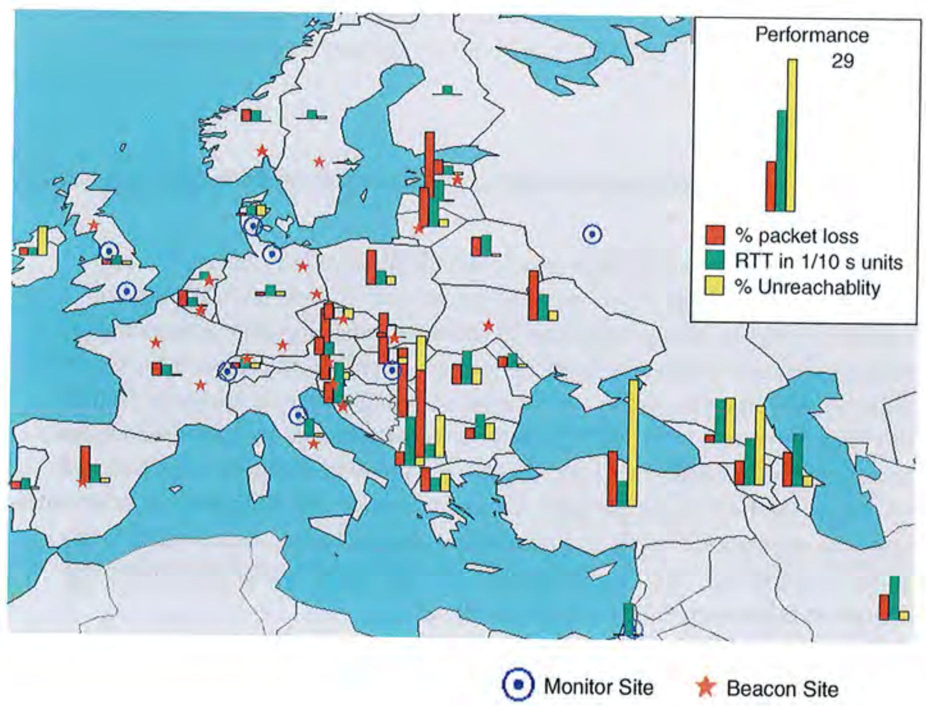

PingER results for end-to-end round-trip times, packet loss, and unreachability for sites in European countries, as seen from SLAC in April 2000 (for other monitoring data, see http://www-iepm.slac.stanford.edu/monitoring/)

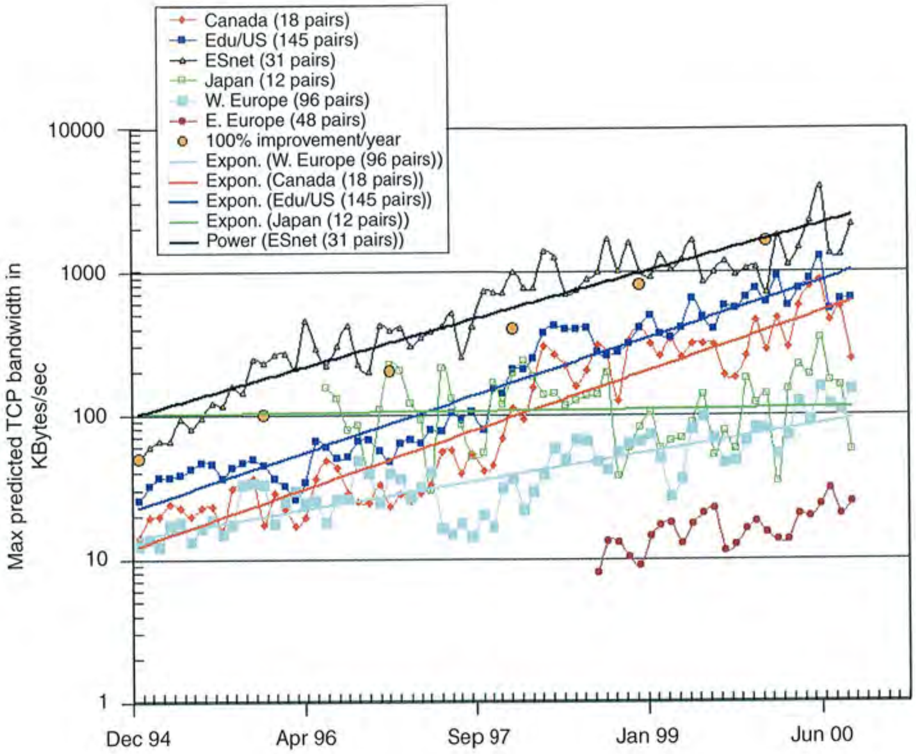

Measured performance quality of groups of sites seen from ESnet monitoring sites from January 1995 through September 2000. The numbers in the legend represent the monitor-site/remote-site pairs contributing to the measurement for each group. The lines are fits to exponentials to facilitate interpretation. The large orange dots indicate a quality improvement of 100% per year. "Quality" is computed as a function of round-trip times and packet loss and corresponds roughly to the maximum TCP bandwidth.

Network Research & Development

Since its inception, ESnet has been a leading-edge network service provider. It has maintained this status by rapidly adopting new services and technologies (e.g., ATM backbone) that meet the needs of its users. The users take advantage of these new network services by modifying existing applications or creating new ones. In addition, as network services improve, more DOE scientific programs move to distributed computing environments (e.g., the Particle Physics Data Grid) and new demands are placed on the network. This process leads to a continuing cycle of network upgrades and application enhancements.

To address this revolving demand cycle, ESnet supports an integrated research and development program, as described on page 46. In this program, application programmers, middleware developers, network engineers, support technicians, and end users work together to tackle the problems encountered in creating large distributed computing environments. Each part of the ESnet community brings a set of assumptions and expertise to bear on the problems it is facing. The other members of the community can verify these assumptions and provide insight into how different proposed solutions would work. This focus on vertical integration allows each part of the DOE community to make significant contributions for all.

Focus areas for network R&D to support ESnet

To maintain its status as a leading-edge provider of network services, ESnet will support three major thrusts in network research and development.

Evaluate and deploy new wide-area and local-area network technologies. *The explosive growth in optical networking technologies is rapidly changing how wide-area networks (WANs) are designed and deployed. The major impact on ESnet will be the replacement of its ATM backbone with new network technologies based on the packet-over-SONET (POS) approach. To accomplish this technology shift, ESnet will experiment with POS and other possible WAN technologies (e.g., multi-protocol label switching [MPLS]). ESnet staff also need to participate in the design and development of new routing and forwarding protocols being developed by the Internet Engineering Task Force (IETF).*

In the local-area network (LAN) environment, rapid changes are also occurring. Standards for 1 Gbps and 10 Gbps Ethernet switching technologies are being developed by the Institute of Electrical and Electronics Engineers (IEEE). While ESnet does not directly support these LANs, its staff must understand the emerging technologies and ensure that WAN capacities match those of the new LANs. ESnet needs to monitor the work being done by IEEE to determine how these emerging standards will affect DOE applications.

In addition, ESnet staff and networking staff at the laboratories must work with vendors to conduct advanced deployment tests of these technologies.

Create and deploy improved versions of current network technologies and services. *Results from recent quality-of-service (QoS) experiments have shown that ESnet routers can be configured to support the emerging IETF differentiated services technology standards. By using this tech-nology, sites are able to indicate the relative priority of a packet without the need for explicit signaling between site routers and ESnet routers. ESnet staff members need to continue to be involved in these experiments to determine when and how to deploy these technologies in production ESnet services.*

Another set of experiments has uncovered serious performance problems that exist in today's network applications and protocols. These problems range from the misconfiguration of network devices (i.e., hosts, routers, and switches) to interactions among multiple protocols. To address these problems, DOE researchers are creating end-to-end measurement tools and components that can be used to determine what specific problems exist. The ESnet testbed provides a unique national backbone infrastructure that allows these tools to gather information on the operation of core network routers. Analysis tools can then be used to correlate events from multiple locations. The results will help users tune applications for best performance and help network managers tune parameters and monitor traffic flowing through network switches and routers. ESnet staff members, working with the users, can monitor these end-to-end experiments to understand how the demands of an application affect the operation of the network.

Conduct horizontal and vertical integration efforts to identify problems in multi-domain and multi-community environments. *In the distributed computing environment of which ESnet is a part, each domain (i.e., lab, university, or backbone network) maintains control of its own network resources (i.e., links, switches, and routers). ESnet staff members are involved in inter-domain working groups (i.e., IETF) to identify potential problems and propose possible solutions. ESnet is not a simple "bit pipe" that blindly moves data from one location to another but rather provides routers and switches that are active components of the environment. Future services (e.g., QoS) will make ESnet an even more complex and active component in future distributed computing environments. To address this complexity, ESnet staff must work with network researchers, middleware and application developers, and end users to understand how applications operate in this complex environment.*

3 Core Program Requirements

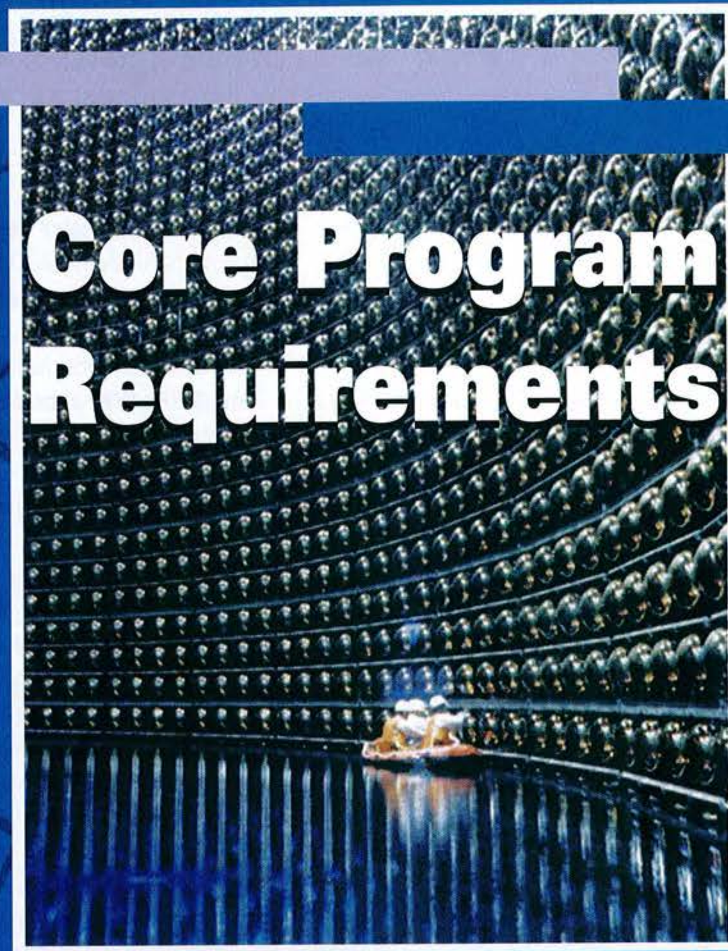

High Energy Physics

ESnet

Program Description

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.

The U.S. HEP program is supported by both 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.

Facilities. Apart from the theoretical component, most investigations are carried out at the major accelerator centers, both in the United States and abroad. The U.S. accelerator laboratories include

- Fermi National Accelerator Laboratory (Fermilab), which operates 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 (SLAC), which operates the new electron-positron "B-factory" (PEP-II), along with a major detector (BaBar) and a fixed-target program at electron energies up to 50 GeV.
- The Cornell Electron Storage Ring (CESR), funded by NSF, which provides electron-positron collisions at up to 10 GeV.

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. Three laboratories that operated HEP accelerator facilities in the past, Argonne National Laboratory (ANL), Lawrence Berkeley National Laboratory (LBNL), and Brookhaven National Laboratory (BNL), 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) that is 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 to 5 GeV.

The ATLAS detector, now under construction for the CERN Large Hadron Collider, has millions of channels of electronics and will observe 100 million proton-proton collisions per second (http://atlasinfo.cern.ch/Atlas/Welcome.html)

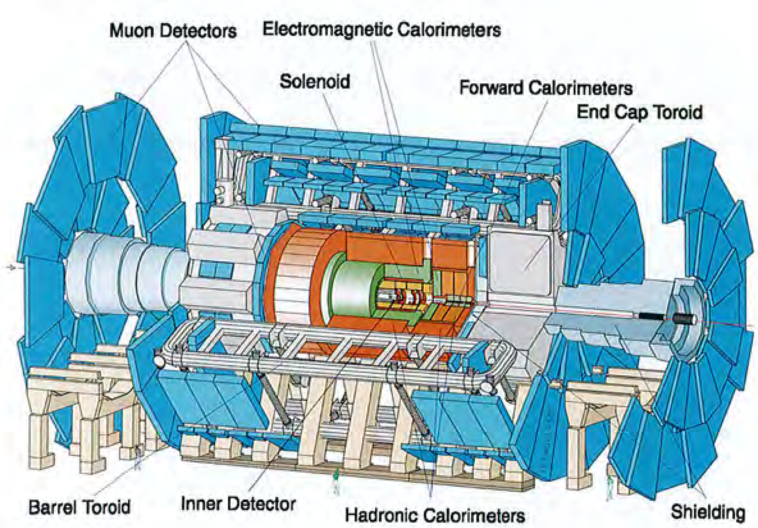

Collaborations. About 140 universities are an integral part of the U.S. HEP program. University physicists work on experiments at accelerator laboratories both in the U.S. and abroad, and on a variety of other experiments. Networking is particularly vital to the HEP program because of the large international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

New programs. HEP researchers from the United States have initiated a major effort at the LHC at CERN. Among other goals, this facility will search for the Higgs boson. Other new aspects of the U.S. HEP program include the PEP-II B-Factory at SLAC, which started taking data late in 1999 with an international community of physicists interested in studying matter-antimatter asymmetry, and the Main Injector and upgraded Tevatron Collider at Fermilab, which will start Run II in 2001, seeking to discover the Higgs particle or related phenomena before the LHC turns on.

Current Program Requirements

Experiments at the major accelerator centers are large-scale enterprises conducted at DOE laboratories and abroad. They typically involve 100 to 500 physicists and at least as many engineers and technicians during the construction phase. These collaborations include projects in which U.S. researchers collaborate in experiments at foreign laboratories, as well as projects in which foreign researchers collaborate in experiments at DOE laboratories in the United States. From initial conception to final data collection, such experiments range in duration from 5 to 15 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 an international mix of institutions.

The primary factor shared by participants in a collaborative HEP experiment is an interest in a certain approach to physics; there is little regard for geographic proximity in forming such collaborations. 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.

The network is also needed to access major facilities, including accelerators, as noted above, and other facilities important to the program. For example, access to supercomputer centers is vital for computationally intensive parts of the program. Theoretical calculations, especially those involving numerical evaluation of field theories on a discrete lattice, rely heavily on access to high-performance computing. Detailed accelerator simulations in which particle orbits are followed in detail are also significant uses of the supercomputers.

Future Program 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, virtual terminals, remote access to files, and basic teleconferencing must operate with complete reliability and at high speeds. The two applications discussed below, while not the only anticipated needs, are expected to be the services most necessary to high energy physics.

Telepresence. In 1990, HEP groups began an experiment in the use of videoconferencing for scientific collaborations. The initial link was established between conference rooms at LBNL and SLAC, and this type of videoconferencing was found to be a highly effective medium for collaborative meetings. In the years since, nearly all groups using ESnet have come to rely on this service and a related service using workstations and wide-area network infrastructure. These two services have formed the basis for the ambitious developments in conferencing and remote operation of facilities that were developed in DOE2000 and related programs (see Chapter 2). HEP experiments and collaborations continue to grow in complexity and will require 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.

Database access. The large data sets expected from Tevatron and PEP-II experiments (100 terabytes per year) in the near future and LHC experiments (petabytes 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 analysis. Most experiment groups envision a system of regional centers that provide local access to the large data sets. The base data sets might not be transmitted over the network (and in any case are not subject to rapid change), but smaller, partially processed "summaries" will be needed. Such summaries can be used for detailed physics analysis by distributed analysis groups and will need to be transmitted and cached in sophisticated ways that anticipate this use.

Connectivity

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

The 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, KEK, and the Chinese Institute for High Energy Physics (IHEP) creates a requirement for higher bandwidth connections to those sites. A major task of the ESnet program to date has been providing 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 for 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, and Internet2 will attempt to provide the improved connectivity so much needed by the university researchers. However, ESnet will have to provide assistance when these efforts do not meet specific needs.

The MINOS experiment will send a beam of neutrinos through the earth from Fermilab in Illinois to the Tower-Soudan mine in northern Minnesota in order to make precise measurements of the very small neutrino masses (http://www.hep.anl.gov/ndk/hypertext/numi.html)

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 among all of the participating institutions. The following are the major networking requirements of the HEP domestic program:

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

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 they are not so today clearly limits the efficiency and productivity of present HEP collaborations. The major requirements for international links during the next five years are shown in the table.

International Connectivity Required for HEP

Destination	Current	5 Years
CERN, Geneva, Switzerland[1]	155 Mbps	1,000 Mbps
DESY, Hamburg, Germany	45	155
Other sites in Europe	45	155
KEK, Tsukuba, Japan (and other sites in Japan)	10	100
IHEP, Beijing, China (and other HEP institutes in China)	1.5	5
Key institutes in Mexico, Brazil, and Argentina[2]	1.5	5
Moscow region in Russia[3]	1	5

1 Steady growth as the LHC nears completion.
2 With connectivity to other parts of Central and South America.
3 Joint Institute for Nuclear Research (JINR) in Dubna, the Institute for High Energy Physics (IHEP) in Protvino, and other points. Connectivity is also needed to institutes in St. Petersburg and Novosibirsk and to institutes and universities throughout Russia and other former Soviet republics.

Aerial photograph of the Fermilab accelerator complex. The large circle houses the Tevatron Collider (http://www.fnal.gov/).

Run II for the Fermilab Tevatron Collider

The highest energy particle accelerator in the world is the Tevatron Collider at Fermi National Accelerator Laboratory. For the past five years, both the Collider and the detectors that turn the collisions into physics data have been undergoing major upgrades. The upgraded collider will provide 150 times as many collisions as were previously observed, and the detectors will have much improved capabilities to identify the particles that emerge from the collisions.

In the Tevatron, beams of protons and antiprotons collide at nearly the speed of light. Seven million times a second, individual protons and antiprotons collide and burst into showers of secondary particles: quarks and electrons, muons, neutrinos, B mesons and W bosons, or any of more than a hundred possibilities. The collisions take place inside each of two huge collider detectors on the accelerator ring. The detectors' job is to observe as many collisions as possible, to recognize and record the resulting particles, and to preserve the information for later study. By analyzing the stored data from the detectors, physicists make discoveries about the fundamental nature of matter and energy.

As with any fundamental research, one cannot know in advance what will be observed, but physicists at the Tevatron believe it can provide evidence for several hypotheses on the list of hot topics in physics. Chief among these are possible explanations for the breaking of electroweak symmetry. In the conventional explanation (given by the "standard model" of particle physics), the effect results from the presence of the still-hypothetical Higgs particle. If that premise is correct, it seems very likely that the Higgs particle could be observed in the next few years at the Tevatron Collider. Another explanation proposed for the effect is the existence of

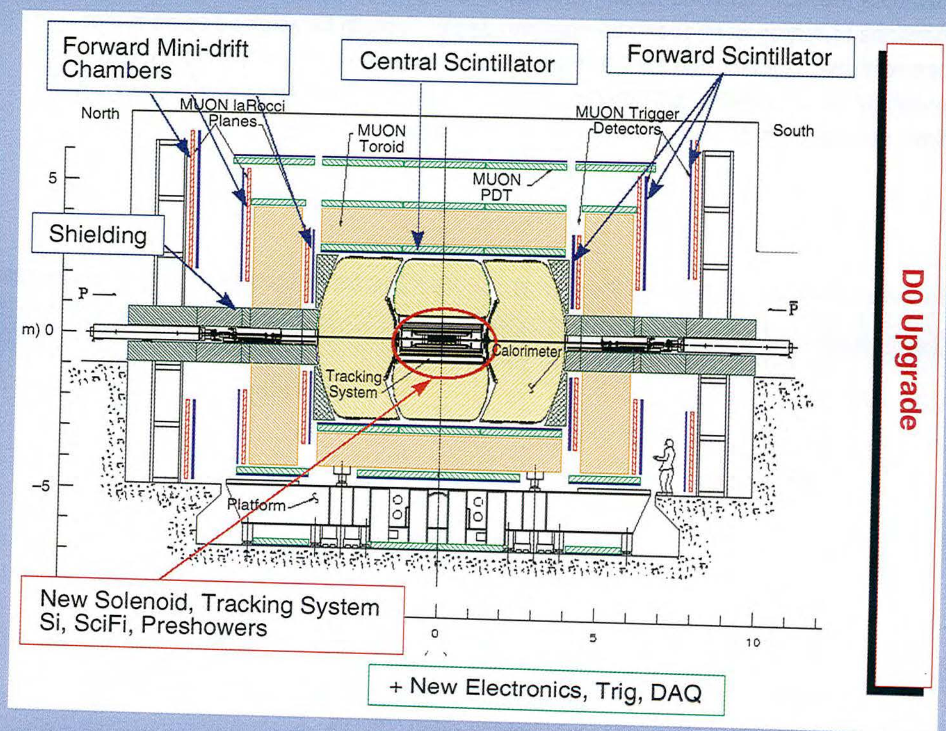

Drawing of the upgraded D0 detector

a set of interactions and corresponding particles referred to as supersymmetry. If this explanation is correct, it also seems likely that evidence to confirm it can be found at the upgraded Collider.

The upgraded detectors, known as CDF and D0, have been built and operated by large international collaborations, each consisting of about 500 physicists from five or more countries on three continents. These collaborations have used ESnet and its partner networks extensively to manage the projects and to share information during the design and construction stages. When the detectors are turned on early in 2001, the collaborations will become an even larger source of traffic on the network as physicists around the world move and access data and share preliminary results in the race to find the Higgs particle...or whatever the answer turns out to be.

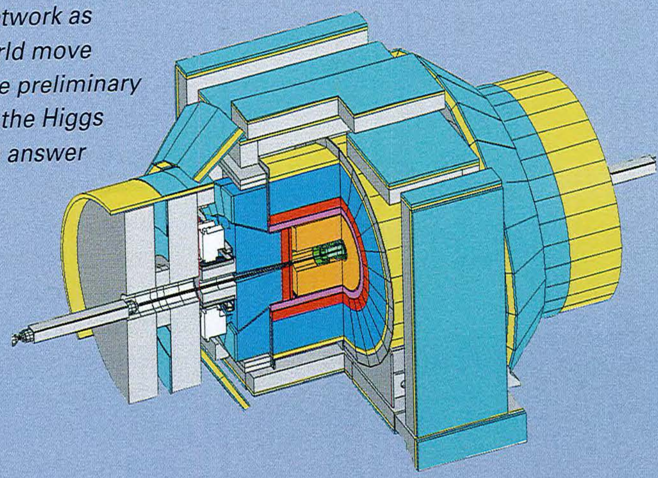

Cutaway view of the upgraded CDF detector

Networking and HEP: Technical Opportunities

Initiatives in which HEP has collaborated with related fields (nuclear physics, particle astrophysics, experimental cosmology) have recently identified virtual data grids (VDGs) as the preferred technical direction for providing access to multi-petabyte data sets by geographically distributed collaborations. Such grids would extend current developments in computational grids. In planning documents prepared for experimental collaborations and in proposals to DOE and NSF, these grids are envisioned as a multi-level hierarchy of distributed servers (both for data and for computation) that provide transparent access to both raw and derived data. They would be used by physicists working to extract answers from experimental data. This vision and the initial proof-of-concept work both depend heavily on high-speed, reliable networks. Since the collaborations are international, increased stress will be placed on smooth connections with international links. One of the challenges of the developing VDGs will be to provide strategies for data caching and delivery when transoceanic links are likely to lag behind the performance of land-based links.

Planning and Management

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 needs sufficient network management resources to prevent interrupted or inadequate service. A crucial role for management is to forecast requirements well enough, based on input from the ESnet Steering Committee and other sources of information on the evolution of DOE programs, to provide improved performance and connectivity by the time they are needed. Such planning should also take into account the delays that may be encountered in deployment once orders are placed. 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.

Super Kamiokande underground detector

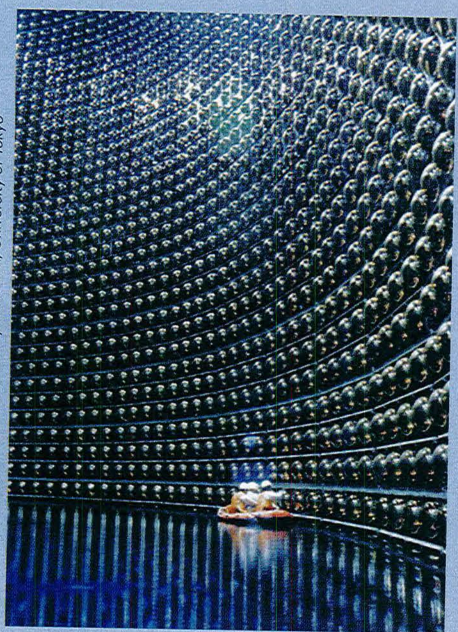

Photo courtesy of Institute for Cosmic Ray Research, University of Tokyo

Photo of the Super-Kamiokande underground detector under construction and drawing of the completed detector with access tunnels (http://www.phys.washington.edu/~superk/)

Researchers using the Super Kamiokande (known as Super-K) detector electrified the physics world during the last two years with their announcement of strong evidence for the phenomenon of neutrino oscillations. Such oscillations could only occur if neutrinos have mass. Previously, in the standard model of particle physics, neutrinos have always been assumed to have no rest mass. The discovery of neutrino mass is interesting not only as new knowledge about somewhat mysterious particles, but also because it could explain why the universe appears to contain at least 10 times more "dark" (unobservable) matter than luminous matter such as stars and galaxies.

The detector consists of an inner volume and an outer volume that contain 32,000 tons and 18,000 tons of pure water, respectively. The outer volume shields the inner detector against entering cosmic ray muons and radiation emitted by the surrounding rock and walls. The inner detector has 11,200 photomultiplier tubes (PMTs) that are attached to the bottom, top, and sides and face inward. The PMTs collect the pale blue light (called Cerenkov light) that is emitted by particles traveling as fast as light through the water. By measuring the direction and intensity of this light, information about particle interactions (such as neutrino interactions or proton decay) can be determined.

The neutrinos from which oscillations have been inferred were produced in the atmosphere by primary cosmic rays. Super-K can also observe neutrinos produced in stars, including those from the sun (which carry information about the nuclear reactions that power the sun) and those created when more distant stars explode as supernovae.

There are four kinds of forces in nature by which particles can interact: the strong, weak, electromagnetic, and gravitational forces. The electroweak theory that unifies the electromagnetic force and the weak force has already been established. The next step of elementary particle physics is to establish the Grand Unified Theories that unify the strong, weak, and electromagnetic forces. The Grand Unified Theories predict that protons, the basic constituent of matter, decay. They also predict the existence of monopoles, elementary particles that weigh as much as bacteria. Researchers are using Super-K to study the Grand Unified Theories experimentally by seeking proton decay and monopoles.

The Super-K detector is located in a working mine at the heart of a mountain in central Japan. It is operated and used for physics studies by an international collaboration with members from four countries. The largest contingents, however, are first from Japan and then from the United States, which is represented by several institutions from across the continental United States and from Hawaii. ESnet and its connections to Japan and Europe are vital to the collaborative process involving Super-K.

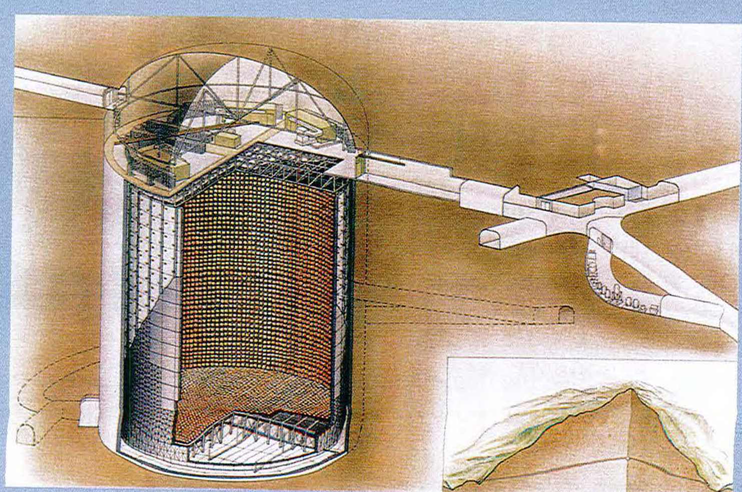

Illustration courtesy of Institute for Cosmic Ray Research, University of Tokyo

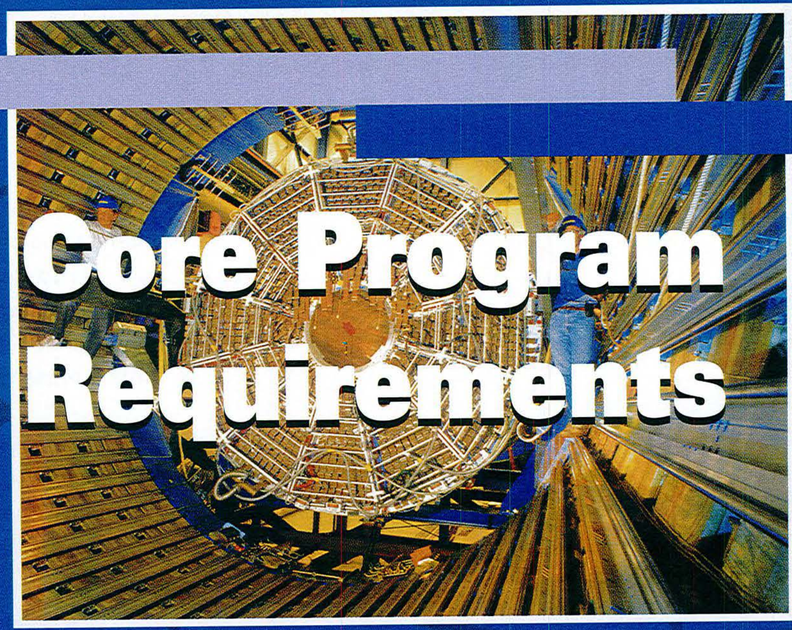

3 Core Program Requirements

Nuclear Physics

ESnet

Program Description

The Nuclear Physics program within the Office of Science seeks to understand the structure and interactions of atomic nuclei, as well as the fundamental forces and particles of nature, as manifested in nuclear matter *(http://www.er.doe.gov/production/henp/np/overview/overview.html)*. It encompasses both large and small experimental facilities for studies in medium energy, heavy ion, and low energy nuclear physics, plus complementary research in nuclear theory.

Medium energy nuclear physics. This program is aimed at understanding the structure of the atomic nucleus in terms of quarks and gluons (the constituents of nuclei), and the strong force (which ties them together). Unraveling the mysteries of their interactions requires a sensitive probe, typically an energetic beam of protons, electrons, or photons.

The most intense such probe is at the Thomas Jefferson National Accelerator Facility (Jefferson Lab; see *http://www.jlab.org/*), which operates a 5.5-GeV electron accelerator capable of simultaneously supporting three experimental halls with a high-current, low-emittance continuous wave beam. Lower-energy machines include the Bates Linear Accelerator at the Massachusetts Institute of Technology (MIT) and the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory. DOE researchers in this field also participate in experiments conducted at the DESY accelerator (Hamburg, Germany).

Heavy ion nuclear physics. Researchers in this program seek to understand the properties and dynamic behavior of atomic nuclei and nuclear matter over the wide range of conditions created in nucleus-nucleus collisions. The program covers a range from low-energy, highly deformed (high spin) nuclei to extremely high energies aimed at observing a predicted form of matter called the quark-gluon plasma.

Two major projects in this area are the experiments known as STAR and PHENIX (see page 55) at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory *(http://www.rhic.bnl.gov/)*. RHIC also hosts two experiments of more modest size, BRAHMS and PHOBOS, which explore specific facets of heavy ion collisions. Finally, the pp2pp experiment will make use of the RHIC accelerator to study polarized proton interactions.

Heavy ions are used at the Argonne Tandem-Linear Accelerator System (ATLAS) facility at Argonne National Laboratory *(http://www.phy.anl.gov/atlas/)* to study highly deformed nuclei. The 88-inch cyclotron at Lawrence Berkeley National Laboratory *(http://www.lbl.gov/)* and smaller machines at Texas A&M University, the University of Washington, and Yale University round out this research program.

Low energy nuclear physics. This program explores the structure of nuclei, nuclear reaction mechanisms, and experimental tests of fundamental symmetries.

The Holifield Radioactive Ion Beam Facility at Oak Ridge National Laboratory *(http://www.ornl.gov/)* provides experimental data as input to refined astrophysical calculations. U.S. researchers are also partnering with Canadian scientists at the Sudbury Neutrino Observatory, where measurements of solar neutrino fluxes relevant to the question of whether neutrinos have mass are anticipated in FY 2001.

Nuclear Theory. The nuclear theory program seeks to provide insight into the observed behavior of atomic nuclei, under a wide range of conditions, particularly those covered by the experimental programs described above. Nuclear theorists from within the national laboratories and many universities collaborate in tackling fundamental issues in the understanding of matter. One emerging trend in nuclear theory is the use of large computers to compute, from first principles, fundamental properties of hadrons (a class of subatomic particles that includes protons and neutrons). This type of calculation, based on lattice quantum chromodynamics, requires enormous computing power. Researchers in this area have proposed three large machines optimized for lattice computing at Brookhaven National Laboratory, Fermi National Accelerator Laboratory, and Jefferson Lab to serve this national user community. ESnet will link researchers to these computing facilities and to parallel efforts in Europe and Japan.

Collider recreates first instants in the life of the universe

In an effort to recreate the conditions of the early universe to gain insights into the fundamental nature of matter, nuclear physicists have recently begun studying what happens when a beam of gold atoms collides at high energy with a gold target. These collisions are created in the experiments known as STAR and PHENIX at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory.

By creating extraordinarily high energy nuclear collisions, the collider will act as a giant pressure cooker, producing temperatures and densities tens of thousands of times greater than exist now even at the center of stars. These extreme conditions should, for a fleeting moment, allow the quarks and gluons within the gold nuclei to exist "freely" in a soup-like plasma, a state of matter that is believed to have last existed millionths of a second after the Big Bang, when the universe was first formed.

Each of the millions of events recorded by the STAR detector is stored for later analysis at the RHIC Computing Facility, which serves the experiments known as STAR, PHENIX, BRAHMS, and PHOBOS. Each of these four experiments is supported by large international collaborations, for whom ESnet is an invaluable tool. For example, the STAR collaboration involves more than 400 scientists and engineers from 33 institutions in 7 countries.

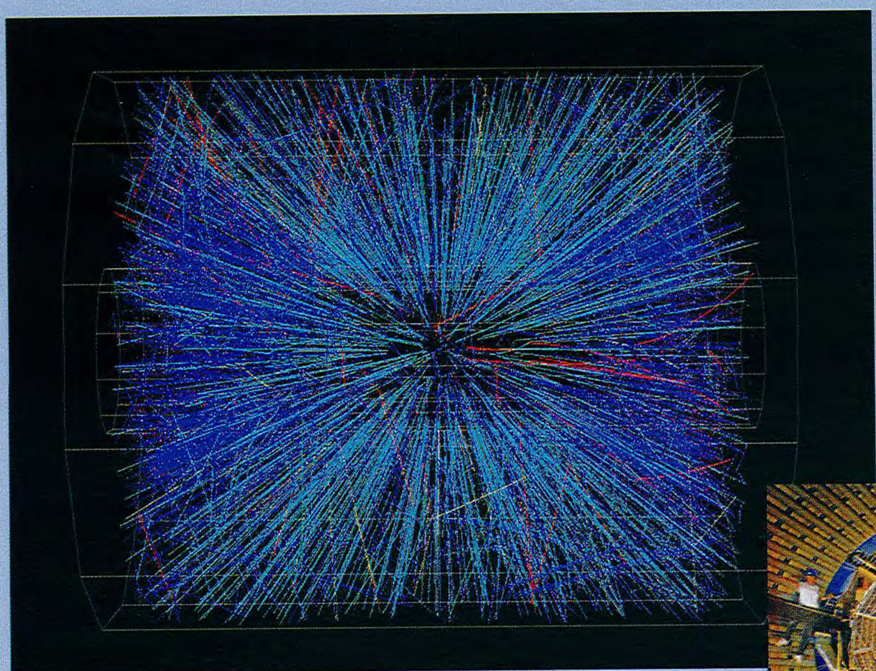

More than a thousand tracks can be seen emerging from one head-on collision between gold nuclei

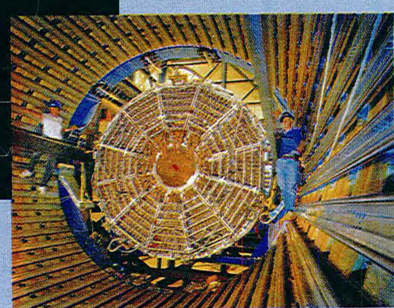

Inside view of the STAR detector

Current Program Requirements

Researchers from industry, universities, and other federal agencies collaborate to make use of the capabilities provided by DOE. In much of this research, DOE partners with the National Science Foundation (NSF) Nuclear Physics program. DOE provides a large portion of the accelerator facilities, and NSF provides a larger percentage of the university researchers. (NSF supports two nuclear physics facilities: the Indiana University Cyclotron Facility and the National Superconducting Cyclotron Laboratory at Michigan State University.) ESnet provides the essential resource that brings these people and facilities together.

Research today typically involves collaborations spanning many institutions and many countries. For example, 1,400 scientists from around the world participate in the research program at Jefferson Lab, and a comparable number are involved in the four RHIC experiments. A major challenge for these collaborations is to keep everyone adequately informed about progress to be exploited and problems to be addressed. A key enabler for this interaction is the Internet, with ESnet providing connectivity for the DOE facilities mentioned above.

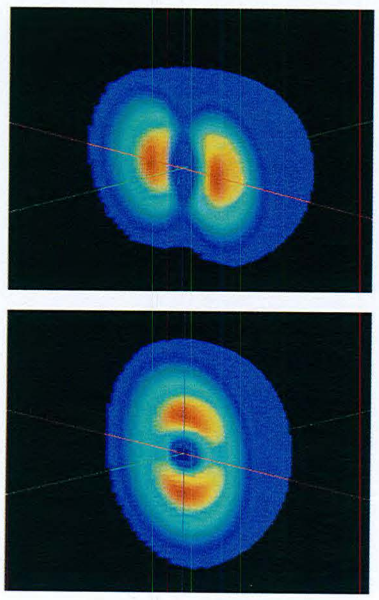

Image of a deuterium nucleus (proton and neutron) in two quantum states, obtained using an electron beam at Jefferson Lab. Such images are used to test theories about how particle interactions affect the shape of a nucleus.

E-mail and simple Web services are the two most widely used communication tools within a collaboration. Their asynchronous nature (author and reader do not have to be involved simultaneously) is particularly important in spanning multiple time zones and busy schedules. Whereas in the past e-mail consisted of short messages, today large design documents, proposals, and research papers are either sent as e-mail attachments or posted to a Web site (or both). This method of communication allows detailed information to rapidly reach an entire worldwide collaboration. As bandwidth has gone up, figures, pictures, and even movie clips have become a valuable part of this communication channel.

Equally mature and valuable are remote login and file transfer capabilities, which support remote access to data and computing capacity. In the past, most interactive use of remote systems was confined to a simple terminal emulation. Today, interactive graphics programs using X-windows, Java, and other technologies enhance the productive use of remote computers and research data. More significant, from a bandwidth point of view, is the need to transfer over the network ever larger data sets, whether experimental data, detector simulations, or theoretical model calculations.

Like the recent generation of high energy physics experiments, the CLAS detector at Jefferson Lab and the STAR and PHENIX detectors at RHIC are each capable of generating more than 100 terabytes of data per year. Even though only a small sample of that data is transferred over the network, an OC-3 (155 Mbps) connection for each of these labs is essential for productive collaborative work.

A case in point is the movement of simulation data to the RHIC computing facility from the simulation generation sites (NERSC at LBNL, the Laboratory for Nuclear Science at MIT, the Pittsburgh Super Computer Center, the RIKEN Computing Center in Japan), and the movement of associated analysis samples out to collaborating institutions. Until recently, such transfers were limited by Brookhaven's T3 connection to ESnet.

While use of international links is not as significant as it is for high energy physics, such connectivity plays an important role in the experiments at both RHIC and Jefferson Lab. These leading-edge research facilities host collaborators from all over the world. International connections through ESnet are key to keeping these collaborations healthy and productive.

Future Program Requirements

Bandwidth

The most significant near-term change in network requirements will be a need for increased bandwidth, particularly in support of the RHIC experimental program. Each of the large detectors there (STAR and PHENIX) is capable of writing 20 megabytes per second to tape, generating several hundred terabytes of data each per year. A single event from a gold-gold collision can generate more than 10 MB of data. Once RHIC reaches full luminosity (before the end of 2001), its data production rate and corresponding analysis activity will saturate the existing OC-3 connection, dictating an upgrade to OC-12 (622 Mbps).

Jefferson Lab experiments will also be slowly increasing their data production rate, but these projects are likely to be adequately served in the near term by the site's recent upgrade to OC-3. In a longer time frame (several years), Jefferson Lab is planning for increased energy and a fourth experimental area, at which point it, too, is likely to need an upgrade to OC-12.

Nuclear Physics Data Grid

As Jefferson Lab deploys the tools coming from the Particle Physics Data Grid project, it expects to see a greater use of network bandwidth (and a corresponding increase in the rate at which physics results are achieved). This grid software will accelerate file transfers and allow cached copies of important files to exist at a number of sites.

Brookhaven National Laboratory, home to the Center for Data-Intensive Computing, is working to apply many of these same capabilities to RHIC-related experiments, to the future Atlas experiment at CERN's Large Hadron Collider, and to a number of other data-intensive science programs at Brookhaven that require access to wide-area networks.

Lattice Meta-Facility

Jefferson Lab and MIT are leading the emergence of a Lattice Hadron Physics Initiative. This collaboration will tie together computing resources at Jefferson Lab, MIT, and several other institutions. Its goal is to model quark-gluon interactions by solving the equations of lattice quantum chromodynamics to compare with experiment. These distributed computing resources can be thought of as a meta-facility to which computationally challenging jobs can be submitted (to execute at one of the participating centers). Simulated data will be gathered at Jefferson Lab and distributed for additional analysis to a number of collaborating sites, again using ESnet and the tools from the Particle Physics Data Grid to increase the productivity of far-flung researchers.

Virtual Shifts

The complex physics detectors at DOE facilities are currently operated by a mixture of laboratory staff and visiting scientists, postdocs, and students. This requirement to be physically present at the laboratory has many side benefits, but it often prevents the involvement of certain specialists at key points in the program. With the emergence of secure authentication of quality-of-service capabilities within ESnet, it will be possible to "stand shift" remotely and remotely assist in diagnosing problems.

While this capability is attractive to researchers at RHIC and Jefferson Lab, granting remote access to a multi-million-dollar detector will require a high level of trust and authentication. Productive (and safe) interaction with the control of these devices will also require certain guarantees on network bandwidth. An investment in these capabilities by ESnet will yield improved efficiency for these valuable facilities and will lower travel costs for researchers.

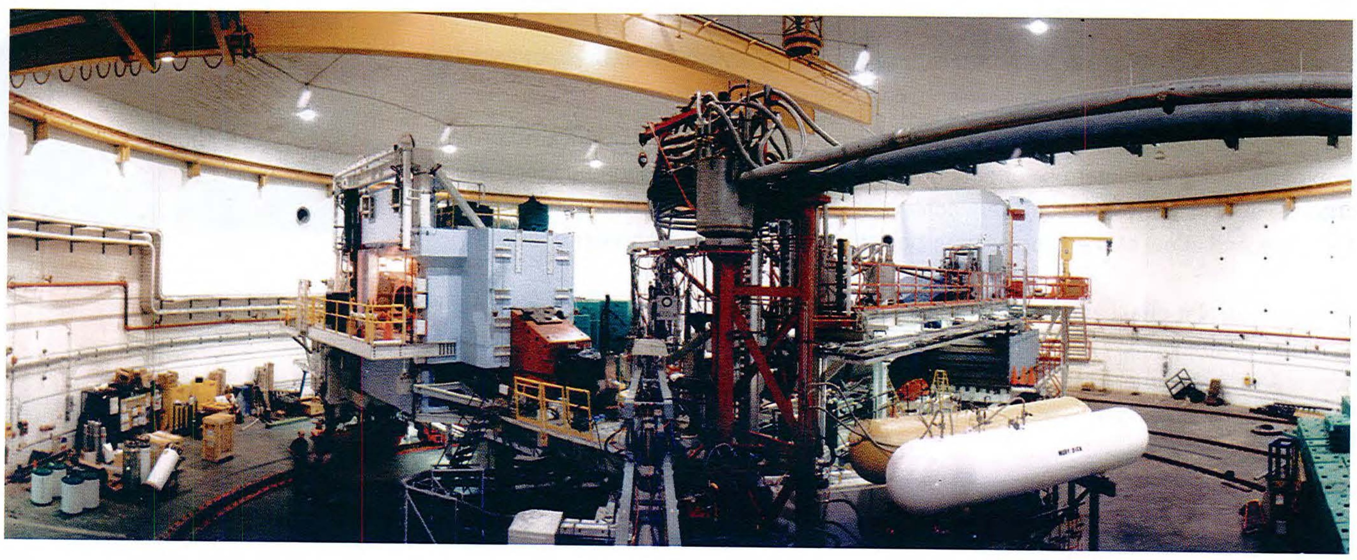

In these two massive detectors at Jefferson Lab, interactions between electron beams and low-mass nuclei reveal the subatomic structure of matter

3 Core Program Requirements

Fusion Energy Sciences

ESnet

Program Description

The long-term mission of the Fusion Energy Sciences (FES) program is to develop nuclear fusion as an environmentally attractive, commercially viable, and sustainable energy source. In the nearer term, the program's goal is to provide the scientific base for the future development of fusion as a power source. To achieve this goal, it has been necessary to develop the science of plasma physics, a field with close links to fluid mechanics, electromagnetism, and nonequilibrium statistical mechanics.

Net power production requires the confinement of hot, dense, ionized gases (plasmas). The intrinsically nonlinear nature of the problem and the abundant free energy available in confined plasmas results in a physical system with complex dynamics and strong turbulence. Combined with the complicated geometry of fusion experiments, these 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 multi-dimensional signals that describe plasma parameters as functions of phase space and time.

Fusion experiments have a dual role: providing data for the advancement of plasma science and increasing plasma parameters (densities, temperatures, discharge times, etc.) toward the levels that would be needed in a power-producing reactor. Progress in FES research is linked to a handful of major experiments that are too expensive to duplicate. As these experiments have increased in size and complexity, there has been concurrent growth in the number and importance of collaborations between large groups at the experimental sites and smaller groups located at universities, industry sites, and national laboratories.

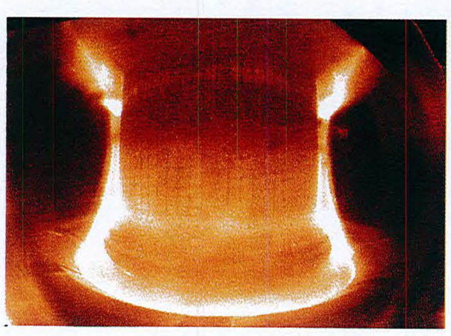

An image of the C-Mod tokamak plasma taken in visible light. Since the plasma temperature is almost 50 million degrees, it emits mainly in the X-ray part of the electromagnetic spectrum. This visible image emphasizes the cooler edge regions of the plasma where it interacts strongly with the walls of the machine. Remote collaborators routinely access such images and associated data via ESnet.

The current generation of experiments in the United States are explicitly "national facilities," each with an advisory committee made up of experts from outside laboratories and universities and a host of collaborators participating in their experimental programs. Three such experiments are the National Spherical Torus Experiment *(http://nstx.pppl.gov)*, the DIII-D facility *(http://fusion.gat.com/diii-d/)*, and the Alcator C-Mod Tokamak Fusion Research Project *(http://www.psfc.mit.edu/cmod)*.

U.S. researchers also work closely with foreign groups and participate actively in experiments in Japan and Western Europe. Various proposals for a "next step" experiment have been put forward. These involve creating so-called "burning" plasmas in which the self-heating power from fusion reactions provides the necessary energy to sustain the reaction. These experiments would certainly be international in scope. Supporting this next step are other facilities used to study the physics of long-pulse (essentially steady state) plasmas.

In addition to experimental programs, DOE's Office of Fusion Energy Sciences supports a large theory and computation effort carried out by staff from national laboratories, universities, and private industry. This effort ranges from the very applied analysis of experimental data to much more fundamental theory and discovery through computational science. Researchers supported by this office make up one of the largest user groups at the National Energy Research Scientific Computing Center (NERSC).

There are two major high-end computing activities:

- Microscopic modeling, that is, studies of turbulence and transport (formerly called the numerical tokamak project).
- Macroscopic modeling, sometimes called magnetohydrodynamics or MHD (see, for example, *http://www.nimrodteam.org*).

In addition to these areas of high-end computer use, there are other significant projects in

- Wave-plasma interaction, which includes energetic-particle physics and radio-frequency-wave heating and current-drive physics.
- Multi-phase interfaces, which includes plasma boundary physics and plasma/wall interactions.
- Design physics (see *http://w3.pppl.gov/topdac/*).
- Inertial fusion energy modeling.

Current Program Requirements

Connectivity

FES research is a worldwide effort conducted at some 90 sites in the United States, 60 in Europe, 40 in Japan, and several in each of South America, China, and Australia. For the U.S. sites, the highest priority for network connectivity is access to computing resources (principally NERSC) and to the major experimental sites at General Atomics, the Massachusetts Institute of Technology (MIT), and the 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.

Fusion energy research has relied on international collaboration since its worldwide declassification in the late 1950s *(http://www.iaea.org/worldatom/inforesource/bulletin/bull374/dolan.html)*. Such collaborations are supported by a wide interchange of ideas and scientific personnel worldwide. In addition, technology is readily exchanged among the countries involved, and components of experimental devices are fabricated all over the world. In the future, we expect that international data links will be as important to FES as our domestic connections are today. Principal foreign sites include the following (for full description, see the list of acronyms on page 110):

- IPP (Institute for Plasma Physics), Max Planck Institute, Garching, Germany.
- JET (Joint European Torus) Joint Undertaking, Abingdon, U.K.
- JAERI (Japanese Atomic Energy Research Institute), Naka, Japan.
- NIFS (National Institute for Fusion Studies), Toki, Japan
- UKAEA Fusion, Culham, U.K.
- DRFC, CEA, Cadarache, France.
- CRPP-EPFL (Ecole Polytechnique Fédérale de Lausanne), Lausanne, Switzerland.
- Korea Basic Science Institute, Taejon, South Korea.
- Institute of Plasma Physics, Max Planck Institute, Greifswald, Germany.
- ENEA Frascati, Italy.

Capacity

Today's ESnet infrastructure, with its OC-3/OC-12 (155/622 Mbps) backbone and its gateways to the broader Internet, generally meets the current domestic needs of the FES program. We have begun moving complex computing applications (real-time remote data access, remote visualization, distributed computing, remote control of experiments, etc.) from a local-area network into the wide-area network environment. This process requires ever-increasing bandwidth.

In this regard, it should be noted that our sites require either fiber distributed data interface (FDDI) or switched 100/1000-Mbps Ethernet or both for their local infrastructure. Requirements for the wide-area network will soon reach these levels. There is concern about the ability of network software and protocols to work with applications that can utilize this bandwidth end-to-end. Architectures that rely on remote access to data from experiments and advanced simulations require real throughputs of 1 to 10 Mbps *per application*—with many applications running simultaneously.

The most notable unmet need is for high-speed access to non-ESnet sites, particularly universities and international sites. Despite improvements from vBNS and Abilene, connectivity to universities is not fast or reliable enough. And while "lifeline" services like e-mail or telnet do function to these sites, network-intensive applications are generally restricted to connections between ESnet sites and a few universities. Most U.S. universities and foreign sites cannot fully participate in applications that require first-rate connectivity.

Services

Fusion energy research requires all basic network services, including e-mail, file transfer, remote login, remote printing and plotting, directory services, and World Wide Web access. In recent years, however, support for remote experimental and theoretical collaborations has added requirements far beyond these lifeline services.

Computing, visualization, and file management for a given task are already shared over the network. 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 includes remote job entry, remote procedure calls, network file service, distributed editors, distributed databases, and so forth. In addition, FES researchers 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 the future network environment.

An initiative that grew out of the Virtual Lab for Technology (sponsored in 1999 and 2000 by the Office of Fusion Energy Sciences) has provided audiovisual services for many control rooms and conference rooms used by FES programs. The nature of the experimental methods employed in FES research requires close communication among all members of the physics teams responsible for acquiring and analyzing experimental data. While many members of each lab's physics team are present at the experimental site for a significant number of runs, it is unlikely that all of them are present every day. Hence, it is necessary to link the control room to collaborators' home sites with data, voice, and video communications. Audio and video are now routinely broadcast from operating experiments via the Web by means of streaming video technologies. Weekly physics meetings from the major experiments are broadcast via a combination of streaming video and Showstation technology. Because current design activities and theoretical collaborations are increasingly multi-institutional and multinational in scope, those researchers have also been deploying more powerful tools for interpersonal communications, such as videoconferencing and multimedia e-mail.

Magnetic Fusion Collaboratory

In addition to the audiovisual services described above, several software and database activities are underway to create a national magnetic fusion collaboratory. The overall objective is to allow scientists at remote sites to participate as fully in experiments and computational activities as if they were working at a common site. These objectives are to be realized through a collaboratory that consists of four components:

- Remote experimental operations.
- Remote code development.
- Online code and data access.
- Planning and coordination tools.

The goal of the collaboratory is to enhance the effectiveness of remote participants, ease the requirements for frequent travel, increase the number of researchers who can participate in experimental operations and analysis, facilitate code development projects involving several institutions, and improve the interaction between theory and experiment by providing greatly improved access to experimental data and to modeling codes.

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 in several projects.

- Scientists located at the University of Wisconsin made turbulence measurements on the Tokamak Fusion Test Reactor at PPPL by using a remote control room from which they could operate their diagnostics while keeping in close contact with their colleagues in Princeton.
- Lawrence Livermore National Laboratory (LLNL) has assembled a remote control room in Livermore, California, in support of a large, long-term collaboration on the DIII-D tokamak hundreds of miles away in San Diego.
- From that control room in California, 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.

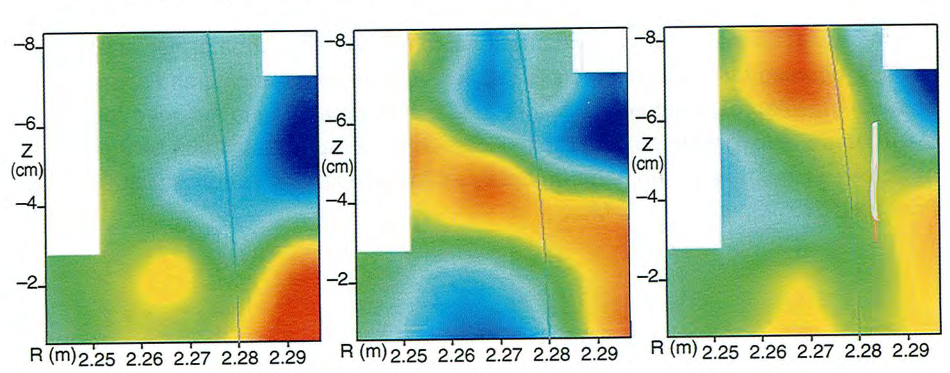

Time- and space-resolved image of density fluctuations in the DIII-D tokamak. Particle and energy confinement in magnetically confined plasmas is dominated by such fluctuations.

These early efforts have been highly successful but are only the first steps needed to demonstrate the technical feasibility of a complete "facilities online" environment.

The major experimental facilities have committed themselves to using the common data formats MDSplus and Microsoft SQL (see *http://www.psfc.mit.edu/mdsplus/*), a shared software library for data access, and a set of shared analysis applications that would use identical data structures and interfaces at every site. A common database structure (schema) is under development that will allow researchers to find runs of every major analysis code. Another database system will allow for translation of variable names between various codes and experimental facilities. The aim is to provide a common environment for all remote and local experimental activities.

Future Program Requirements

Scientific Computing Initiatives

In FY 2000, FES began a new initiative called the Plasma Science Advanced Computing Initiative (PSACI; see *http://w3.pppl.gov/ssp/*). PSACI may become part of the DOE-wide computing initiative called Scientific Discovery through Advanced Computing (SciDAC) if that program receives congressional approval. These activities are based on the formation of community-wide teams to pursue large-scale computing in a collaborative environment. Two teams currently exist, in macroscopic modeling and in turbulence and transport, but more may be formed.

While most of the high-end computing from the PSACI initiatives is performed at NERSC, some is also done at the 256-processor SGI ORIGIN 2000 at Los Alamos National Laboratory and at smaller cluster facilities local to some of the member participants. However, a data-management effort is beginning that involves transferring the output of code runs to disk farms at General Atomics and possibly other locations. This approach will take advantage of the MDSplus storage and retrieval system developed by the FES experimental program. The raw code output needs to be analyzed and visualized, and several high-end visualization facilities are available to the fusion community, including the high-resolution display wall at PPPL and the CAVE facility at the University of California, Los Angeles. The transfer of code output from NERSC to the disk farms to the visualization facility—or to the researcher's desktop visualization screen—will require significant bandwidth.

The increased formation of multi-institutional teams for large code projects also implies the need for improved tools to support remote conferencing and shared software development projects. Another area that is expanding is the need for distributed theory groups to gain rapid access to the experimental data being produced on the three major experimental facilities at General Atomics (DIII-D), Princeton (NSTX), and MIT (C-Mod).

The international collaboration program in the theory and computations area is also growing. Major collaboration programs are in place with the JET facility in England and the JT-60 facility in Japan, as well as smaller collaborations with MAST in England, ASDEX-U in Germany, Tore Supra in France, and KSTAR in Korea. There is increasing need for higher bandwidth connections with these facilities to support data transfer as well as transfer, analysis, and visualization of code output and collaboration activities, including videoconferencing.

Experimental Collaborations

Collaborations among U.S. researchers on domestic experiments will continue to grow. An increasing number of plasma diagnostic systems are routinely operated remotely. Data systems and analysis tools now under development will greatly ease access to large data archives and thus lead to more frequent and larger transfers of data. Experiments are moving more data from "offline" storage systems (magnetic tape and optical media) to "online" magnetic disks. Soon, researchers will probably sift routinely through multi-terabyte archives from remote locations. While architectures that minimize the transfer of such data across the wide-area network will be employed, it is inevitable that the quantity of data moving across these links will increase. Analysis codes will be more closely linked to the experimental data systems. These codes will run on both local and remote systems and require upgraded communications. Universities will play an increasingly important role in these collaborations, putting further stress on an already weak link in the communications system.

At the same time, audiovisual communications requirements will continue to grow. Staff and planning meetings are now routinely broadcast to remote participants. The number of participants at these meetings and the amount of data that will be shared are both expected to increase. Real-time communications during experimental operations will grow significantly, especially one-on-one remote conferencing and the use of display-sharing tools. Operation of fusion experiments presents a unique challenge for collaboration tools. In a typical control room, 50 to 100 scientists and engineers work together essentially in real time. These experiments are not "user facilities" in the ordinary sense. There is a substantial premium for doing as much analysis and visualization as possible during the experiments because operation is generally guided by the results of shots as they are taken. When parts of the team are distributed around the country (or world), the need for powerful and flexible communications tools is multiplied.

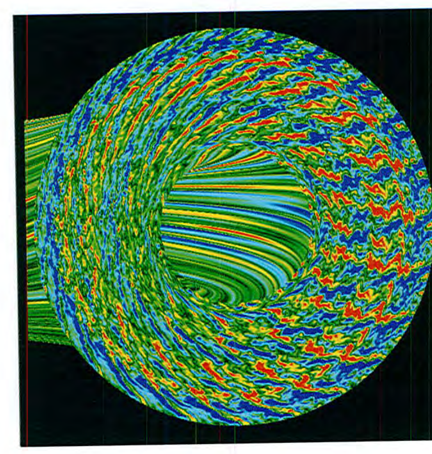

Image courtesy of General Atomics

Gyrokinetic simulation showing contours of the flow pattern of plasma in a tokamak fusion reactor, calculated by the Numerical Tokamak Project, part of the DOE Grand Challenge program

U.S. scientists will be active on a number of international experiments that are currently operating or under construction. These include JET in the United Kingdom, JT-60 and LHD (Large Helical Detector) in Japan, KSTAR in Korea, and Wendelstein 7-AX in Germany. Activities involve scenario modeling; diagnostics subsystems; physics analysis codes; and plasma heating, current drive, or fueling systems. In most cases, U.S. involvement in the international programs will be ongoing and significant. Aggregate traffic levels may not reach those of the domestic collaborations, but instantaneous response, interactivity, and service will have to be greatly improved.

Burning Plasma Experiments

A number of designs for "next step" machines are under consideration by the U.S. and international fusion communities (e.g., see *http://fire.pppl.gov/*). These are principally deuterium-tritium burning devices whose self-generated fusion power would be used to heat and sustain the confined plasmas. The cost and scale of such devices almost guarantees that they will be the product of an international effort. The network demands from the current design activities are quite low. As they go forward, detailed analysis and simulations will begin to have a significant impact. Eventually, if a machine is built, its operation will become a dominant component of the FES program. Researchers anticipate around-the-clock operation by teams located in the United States, Europe, and Japan.

Advanced Computing/ Grand Challenge

As discussed above regarding program activities, FES has launched a new advanced computing initiative, PSACI. The scientific activities in this program (see page 69) will require associated computer science and enabling technology programs for improving data management and visualization.

The advanced computing initiatives will utilize the most sophisticated hardware available, which, in current plans, will continue to be located at NERSC. Typical large-scale simulation runs in this initiative will generate a few terabytes of data, leading to a potential storage requirement of the order of 100 terabytes per year. The program will need not only archival storage of this data, but also a database management system that keeps track of the data, supports searches, and provides a convenient interface for data manipulation and visualization. Given the geographical diversity of the plasma sciences community, any plan to analyze terabytes of data from simulation runs presupposes the availability of a high-speed network. An upper limit for the data transmission rate of a gigabyte per second would be required to transport the data between the production computer and the local visualization station. It is imperative that there be a close coupling between the work undertaken in this initiative and ESnet.

The geographical diversity of the fusion community also requires solid tools for collaborative code development. FES has several such efforts currently underway. Presently the community is wired for low-bandwidth collaborative tools, but with the enhanced computing capability created under the PSACI program, these tools will need to be greatly expanded.

National Transport Code Collaboration

The National Transport Code Collaboration (NTCC) is a project with the goal of changing the way fusion modeling codes are constructed and used in the fusion community. The objective is to develop common, modern transport codes to address major physics issues facing the fusion program. To this end, the project promotes sharing and community ownership of transport code modules and the development of a new framework based on modern computing techniques for rapidly assembling customizable transport codes. Development of national transport codes is based on collaboration among theorists, modelers, and experimentalists located in many of the institutions engaged in fusion research.

This activity is still in its early stages. However, one direction being pursued is the increased use of the network to make experimental data and software tools readily available to plasma physics research groups across the United States and, to a lesser extent, around the world. The network demands of this activity are not envisioned to be as severe as those in the advanced computing/grand challenge area, but the number of sites needing mid-bandwidth connectivity should expand to include all research universities with any plasma physics program *(http://w3.pppl.gov/ntcc)*.

Advanced Visualization Tools

As our computational research moves towards the use of higher dimensional data objects, the role of advanced interactive visualization tools becomes even more important. These tools are needed for understanding the data. Visualization tools are needed not only to display data, but also to find correlations, to visualize subspaces, and to understand complex comparative parametric dependencies. Certain fusion sites now have large high-resolution visual display walls. Effective use of these walls requires not only rapid access to the data, which is normally stored at a remote location, but also a fast local network to allow the parallel computations needed for rapid exploration and rendering of data sets.

Computational challenges in fusion energy sciences

Complete computational modeling of fusion plasmas and the full range of their associated instabilities is well beyond the resources of today's most powerful terascale computing facilities because of the enormous range of time and space scales at play in thermonuclear plasma.

However, many years of close interchange among theoretical analysis, experimental research, and computational modeling have resulted in several reduced (i.e., simplified) computational descriptions of plasmas that are capable of predicting important subsets of the relevant plasma phenomena. While reduced compared to the full first-principles mathematical description of a plasma, the major reduced descriptions themselves present grand-challenge-scale problems, although of a scope that can profitably be addressed on today's computers. This computational effort has several focus areas, corresponding to the leading reduced models of a fusion plasma. These models exist in both the magnetic fusion energy (MFE) and inertial fusion energy (IFE) areas.

Magnetic fusion energy. *In the MFE area, the primary activity is in two categories, those that address the microscopic scales that determine turbulent transport and those that address the macroscopic scales associated with the onset and evolution of rapid, large-scale instabilities. The microturbulence effort has the ultimate goal of developing accurate, physics-based predictive models of anomalous transport in magnetically confined plasmas. The reduced physics model used here is the gyro-kinetic equation, which is a mathematical-physics description of a strongly magnetized plasma that averages analytically over the fast cyclotron motion of the ions and is thus free from the fastest time and spatial scales. This description has proved very useful in calculating the instability driven by the ion temperature*

(continued)

gradient, which is thought to be the dominant microstability causing anomalous transport under normal operating conditions of a fusion experiment.

The macroscopic simulation effort is based on sets of magneto-fluid equations for magnetized plasma that include the effects of realistic geometry and boundary conditions. A goal is to better understand the conditions under which large-scale instabilities develop in hot fusion plasma, as well as the nonlinear evolution and saturation mechanisms. Some of the studies involve the use of simplified geometry, but the overall goal in this area is not merely to study model problems in isolation, but rather to apply predictive simulation capabilities to understand plasmas behavior in actual fusion devices or in nature.

In addition to these major thrust areas, several other large computational efforts in MFE also require substantial computer resources. One is the numerical computation of wave-plasma interactions in realistic geometries. This project is aimed at developing quantitatively accurate predictive models for electromagnetic wave processes that support important heating and current-drive applications in fusion plasmas. Another associated effort is in predicting the behavior of the plasma edge region that connects the hot core plasma to the material surface of the first wall. Interactions in this region create transport barriers that are important to the behavior of the plasma as a whole. Finally, an integrated modeling approach is being used to represent an entire experimental device over macroscopic time scales. All relevant physical processes are represented at some level by interconnected models. The new resources made available by the SciDAC initiatives will enable integrated modeling to occur at a level of realism and complexity much greater than is presently possible.

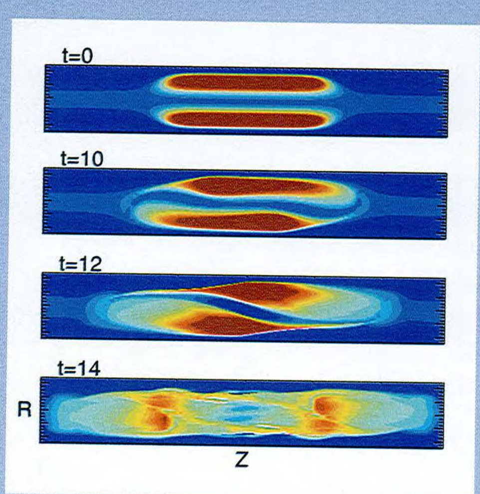

Four frames in a three-dimensional magnetohydrodynamic simulation of a confinement configuration called the field-reversed configuration. Plasma pressure contours are shown. In this simulation, the plasma is described by magneto-fluid equations.

Inertial fusion energy. In IFE, an effort is underway to develop the capability to simulate a "source-to-target" heavy-ion fusion beam. This new capability integrates three types of existing IFE simulations, each of which by itself requires terascale computing: (1) particle-in-cell simulations of acceleration and confinement of the strongly space-charge-dominated ion beams through the driver; (2) electromagnetic and magneto-inductive (Darwin) simulations that describe the beam and the fusion chamber environment and return current effects; and (3) detailed simulations that examine the electron effects and collective modes in the driver and chamber for the final focusing of the beams.

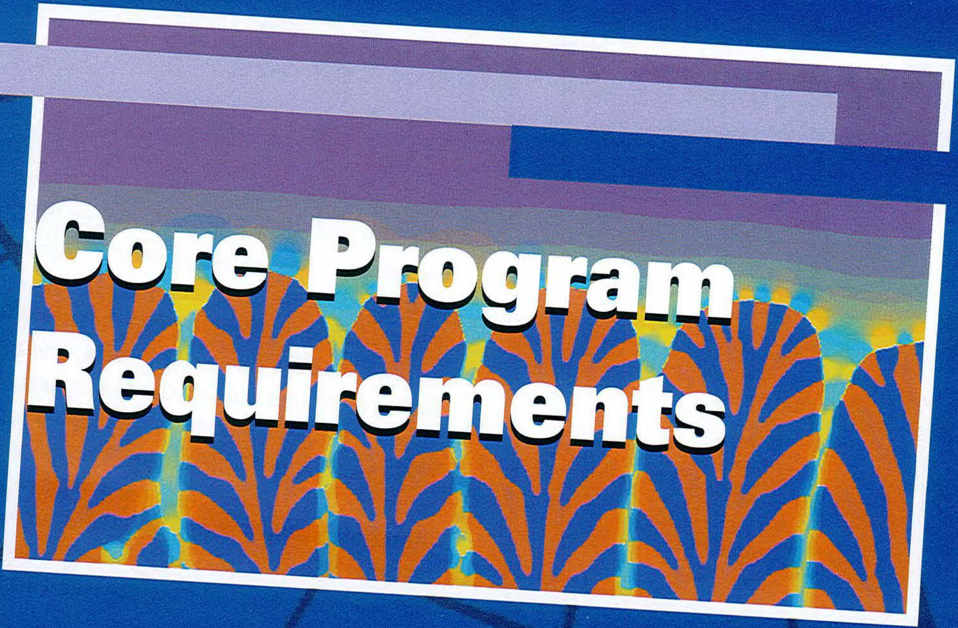

3 Core Program Requirements

Basic Energy Sciences

ESnet

Program Description

The Office of Basic Energy Sciences (BES) supports fundamental research in the natural sciences and engineering to provide a basis for new and improved energy technologies and to provide the understanding necessary for mitigating the environmental impacts of energy use *(http://www.sc.doe.gov/production/bes/)*. As part of its mission, BES plans, constructs, and operates major scientific user facilities to serve researchers at universities, national laboratories, and industrial laboratories. These facilities serve more than 2,400 scientists at 200 U.S. research institutions. BES research is administered through the following subprogram areas.

Materials Science supports basic research in condensed matter physics, metal and ceramic sciences, and materials chemistry. This basic research seeks to understand the atomistic basis of the properties and behavior of materials and how to make materials perform better at acceptable cost through new methods of synthesis and processing *(http://www.sc.doe.gov/production/bes/dms/dmshome.html)*.

Engineering Research conducts fundamental research in support of related technology programs within DOE. The goal of its program is to extend the body of knowledge underlying current energy engineering practice and to broaden the technical and conceptual base for solving future engineering problems in energy technologies *(http://www.sc.doe.gov/production/bes/eng/enghome.html)*.

Chemical Sciences focuses on photo- and radiation chemistry; chemical physics; atomic, molecular, and optical science; catalysis; separation science; heavy element chemistry; and aspects of chemical engineering. This research provides a foundation for understanding fundamental interactions of atoms, molecules, and ions with photons and electrons. This work also underpins a fundamental understanding of chemical reactivity *(http://www.sc.doe.gov/production/bes/chm/chmhome.html)*.

Modeling microstructural pattern formation during the growth of a solid helps identify the final properties of the fabricated material

Geosciences Research supports research aimed at developing an understanding of fundamental earth processes that can be used as a foundation for efficient, effective, and environmentally sound use of energy resources. Such work also provides an improved scientific basis for advanced energy and environmental technologies *(http://www.sc.doe.gov/production/bes/geo/geohome.html)*.

Energy Biosciences supports fundamental research related to a molecular-level understanding of the formation, storage, and interconversion of energy by plants and microorganisms *(http://www.sc.doe.gov/production/bes/eb/ebhome.html)*.

Experimental Program to Stimulate Competitive Research (EPSCoR) is part of BES and supports basic research spanning the entire range of research supported by DOE in states that have historically received relatively less federal research funding. The EPSCoR program supports research cluster activities at EPSCoR states through implementation grants and individual investigator projects through laboratory-partnership grants in more than 15 states and the Commonwealth of Puerto Rico *(http://www.sc.doe.gov/production/bes/EPSCoR/index.htm)*.

User facilities operated by the Office of Basic Energy Sciences

The diverse research programs within BES often require advanced research instruments and facilities that are not commercially available. Such resources are needed, for example, to characterize the structure, properties, behavior, and performance of materials from the macroscopic to the nanometer level. To meet this need, BES operates 18 national user facilities and collaborative research centers (http://www.sc.doe.gov/production/bes/BESfacilities.htm) that are used by its scientists and, when appropriate, by researchers in other agencies, academia, and industry. These facilities represent a major national resource that could not be duplicated economically elsewhere, and often they yield new knowledge that could not be obtained by any other means. Each year, more than 7,000 scientists conduct experiments at the BES user facilities. Thousands of other researchers collaborate with these users to analyze the collected data and publish their findings in peer-reviewed journals. Each of these BES programs and user facilities has network service requirements of its own, as well as requirements common to all.

Synchrotron Radiation Light Sources

- Advanced Light Source (ALS)
  Lawrence Berkeley National Laboratory
- Advanced Photon Source (APS)
  Argonne National Laboratory
- National Synchrotron Light Source (NSLS)
  Brookhaven National Laboratory
- Stanford Synchrotron Radiation Laboratory (SSRL)
  Stanford Linear Accelerator Center

High-Flux Neutron Sources

- High Flux Isotope Reactor (HFIR)
  Oak Ridge National Laboratory
- High Flux Beam Reactor (HFBR) *(closed)*
  Brookhaven National Laboratory
- Intense Pulsed Neutron Source (IPNS)
  Argonne National Laboratory
- Los Alamos Neutron Science Center (LANSCE)
  Los Alamos National Laboratory
- Radiochemical Engineering Development Center (REDC)
  Oak Ridge National Laboratory
- Spallation Neutron Source (SNS) *(under construction)*
  Oak Ridge National Laboratory

Electron Beam Microcharacterization Centers

- Center for Microanalysis of Materials (CMM)
  University of Illinois
- Electron Microscopy Center (EMC)
  Argonne National Laboratory
- National Center for Electron Microscopy (NCEM)
  Lawrence Berkeley National Laboratory
- Shared Research Equipment (SHaRE) Program
  Oak Ridge National Laboratory

Specialized Single-Purpose Centers

- Combustion Research Facility (CRF)
  Sandia National Laboratories
- James R. Macdonald Laboratory
  Kansas State University
- Materials Preparation Center (MPC)
  Ames Laboratory
- Pulse Radiolysis Facility
  University of Notre Dame
- Surface Modification and Characterization (SMAC) Research Center
  Oak Ridge National Laboratory

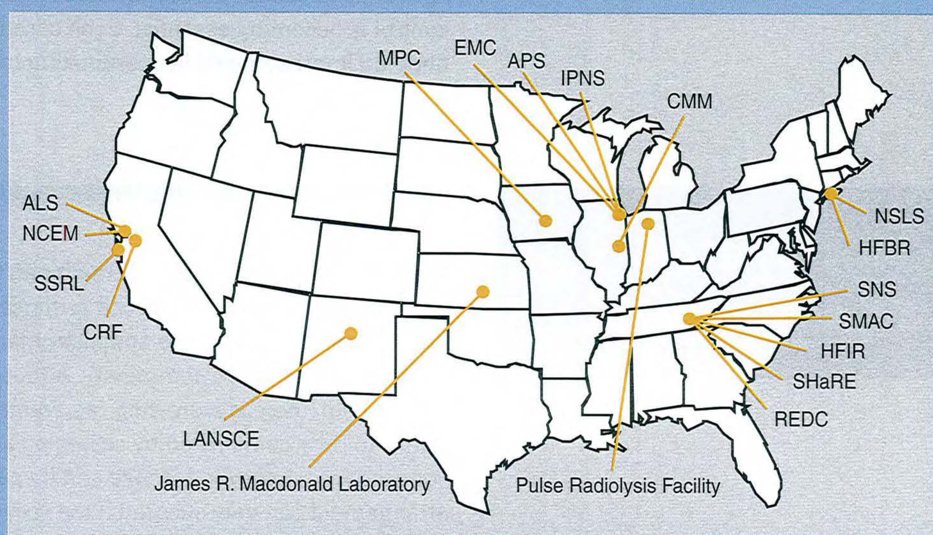

More information on all of these facilities is available through the BES Web site (http://www.sc.doe.gov/production/bes/BESfacilities.htm)

Current Program Requirements

General Requirements

More and more frequently, DOE research programs require the integration of multidisciplinary efforts. In addition, BES scientists are increasingly likely to use experimental and computational facilities across the DOE research complex, which spans the continent. Programs often involve both a research group within a site and collaborating groups at other laboratories, national facilities, or universities. With this in mind, the BES program requirements for networked computing involve nonlocal access to

- Scientific computing resources.
- Data archives.
- Collaborators.
- Experimental facilities.

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 have grown, network requirements have also increased. Similarly, researchers are now able to use remote computational resources to model a phenomenon, but they must also be able to receive the results of these calculations in a timely manner, and such data sets can become enormous.

Connectivity to (and among) the DOE national laboratories and other user facility locations is paramount for BES. This requirement includes connectivity to the widespread set of universities that use DOE facilities. It is also recognized that not all components are within the purview of ESnet, and thus it is key for ESnet to work with other providers of network services to achieve the needed level of connectivity. Many research activities today are designed to heavily exploit the connectivity of the network and cannot proceed without reliable connections, and quality-of-service control is becoming essential. Such control will require both new services and ongoing improvements in coordinating trouble-tracking among ESnet and its peers.

Collaboratory Requirements

Electronic collaboration is becoming part of day-to-day life in cutting-edge science, particularly at user facilities such as those operated by BES. These facilities are unique centers of expertise and instruments for not only the national but also the international scientific community. Although many scientists travel to these facilities to use them, there is a growing requirement to reduce travel and increase efficiency. Both can be accomplished by accessing and operating these resources via network connections whenever possible and by participating remotely with onsite users. This arrangement can only succeed if the access, capabilities, and information available to remote users become as good as if the user were on site. Improvements in network infrastructure, bandwidth, and quality of service are paramount to insuring that these new modes of operation continue to grow.

The broad scope of BES programs, the geographic distribution of BES resources, and the heterogeneous nature of computing equipment pose significant challenges for creating and deploying the needed capabilities. Crucial issues are interoperability (the ability of applications on different kinds of computers to work together) and extensibility (the ability to extend applications with new functions).

For BES 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, but that investigator's presence must be felt at the facility site. The emerging paradigm makes heavy use of videoconferencing, shared screens and shared visualization, remote operation, and electronic notebooks, all of which require much more 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.

BES facilities have participated in several successful pilot projects in electronic collaboration. Remote monitoring or operation of BES experimental facilities has had excellent initial success. Two such projects are the Materials Microcharacterization Collaboratory, which concentrates on experimental resources (see page 76) and the Diesel Combustion Collaboratory (jointly funded with DOE's Office of Energy Efficiency and Renewable Energy and Office of Defense Programs), which focuses on remote access to computational resources and archival data. Both rely heavily on network-based videoconferencing resources to connect remote users and local experts.

Materials Microcharacterization Collaboratory (MMC). This project has integrated and coordinated the efforts of five national centers of excellence in materials microcharacterization across the country into a virtual laboratory for study of materials on scales ranging from tenths of a millimeter to tenths of a nanometer *(http://tpm.amc.anl.gov/mmc/).* In addition to the five laboratories, several instrument companies are partners in the collaboratory.

Networked "collaboratories" allow experts around the world to operate unique instruments, observe data as it is recorded, and consult with colleagues during the experiment—all without leaving home

Materials Microcharacterization Collaboratory

The partner institutions in the Materials Microcharacterization Collaboratory (MMC) collectively house virtually every characterization technique that employs electrons, ions, photons, X-rays, neutrons, mechanical, or electromagnetic radiation to elucidate the microstructure of matter.

Fostering creative exchange. *In the MMC, these centers of excellence are joined into a single, online, interactive, virtual laboratory. In it, creative scientists, educators, and even students are given access to a new environment, one in which convenient, rapid, and dynamic interactions flow unencumbered by the limits of time, space, and location. Participants find that this kind of interchange enhances their ability to conceive and execute new science.*

Users are able to log in and remotely participate in online dynamic experiments and, with proper authorization, control the research tools from their desktop computers via the network. Portions of the collaboratory architecture developed by the MMC are now being used in instruments being developed by the industrial partners in the project, as well as in other research laboratories around the world.

Leveraging resources. *The instrumentation in these centers represents an investment in state-of-the-art technology of over $50 million. The associated research staff—a total user base of more than 500 scientists, students, and industrial researchers—has more than 1,000 man-years of expertise in materials research and characterization. The electronic virtual environment and the interactive tools needed for telepresence microscopy and microanalysis span the breadth of applications and problems encountered by typical materials scientists today. The MMC, therefore, represents a focused implementation of telepresence technologies to scientific instrumentation and research, and the expansion of the MMC to include a wide array of materials science facilities is a natural and anticipated extension for BES.*

Collaboration technologies allow members of the MMC to interact, share data, and control instruments in an online environment (see http://tpm.amc.anl.gov/mmc/). For example, they can use the Web to view colleagues in several control rooms simultaneously, as shown.

Diesel Combustion Collaboratory (DCC). This project is investigating the use of computer and networking technologies to improve communication among participants in the Heavy Duty Diesel Combustion CRADA (Cooperative Research and Development Agreement) *(http://www-collab.ca.sandia.gov/)*. Partners include several national laboratories, a university, and manufacturers of diesel engines.

The goal of the DCC is to make it easier to *do* science, that is, to make the *process* of science, engineering, and information exchange among the CRADA partners more efficient. For example, researchers are able to tackle new problems with their existing techniques simply because these techniques are now easier to use. This project leverages several other efforts to build more realistic physical models of combustion and to design more powerful diagnostic experiments. The DCC has concentrated on creating a repository that allows collaboratory members to archive and share experimental data, model data, and presentation material by using various Web-based tools and to run modeling codes at remote sites from the desktop.

Future Program Requirements

The programmatic diversity within BES is reflected in the range of computational and network resources required by BES investigators and in the use of ESnet facilities and services by these investigators. Computation, information processing, and electronic communications are taking on particular importance in the fundamental research programs of BES that will continue over the next decade. A parallel development will be an increasing need to distribute and share data, expertise, and experimental resources with the BES scientific community. The facilities provided by ESnet will be essential in this role. In addition, three new initiatives will become major users of the network in the coming years, namely, the Spallation Neutron Source, the Nanoscience Initiative, and the Computational Materials Science Network.

Spallation Neutron Source

The purpose of the Spallation Neutron Source (SNS) Project is to provide a next-generation short-pulse spallation neutron source for neutron scattering. The SNS will be used by researchers from academia, national laboratories, and industry for basic and applied research and for technology development in the fields of condensed matter physics, materials science, magnetic materials, polymers and complex fluids, chemistry, biology, and engineering. It is anticipated that the facility will be used by 1,000 to 2,000 scientists and engineers annually and that it will meet the nation's need for neutron science capabilities well into the 21st century.

When completed in 2006, the SNS will be more than 10 times as powerful as the best spallation neutron source now in existence. The experimental data and resources of this will be shared with scientists worldwide using ESnet resources.

Nanoscience Initiative

BES has been a leader in the early development of nanoscale science, engineering, and technology since the 1980s, supporting research and sponsoring workshops to help establish the importance of nanostructured materials. Because of the confluence of advances in this area during the past decade, BES is proposing a major effort to take advantage of opportunities now emerging.

This effort has the following broad goals: (1) to attain a fundamental scientific understanding of nanoscale phenomena, particularly collective phenomena; (2) to achieve the ability to design and synthesize materials at the atomic level to produce materials with desired properties and functions; (3) to attain a fundamental understanding of the processes by which living organisms create materials and functional complexes to serve as a guide and a benchmark by which to measure our progress in synthetic design and synthesis; and (4) to develop experimental characterization tools and theory/modeling/simulation tools necessary to drive the nanoscale revolution. Because of the broad scope of nanoscience, large numbers of scientists will be working collaboratively, in large measure through electronic environments supported by ESnet.

Computational Materials Science Network

The tremendous progress of computational materials science is creating exciting new areas of research, many of which involve the study of highly complex materials or phenomena. In such work, progress is more rapid when researchers work in teams. With these ideas in mind, BES is in starting the Computational Materials Sciences Network (CMSN). This network will link diverse teams of researchers who are committed to working together to solve outstanding materials problems that require cooperation across organizational and disciplinary boundaries.

These collaborative research teams in the CMSN will be developing and using parallel computer codes (on various BES computing resources such as the Argonne IBM-SP, the Cray T3E at NERSC, and BES-funded Beowulf clusters). They will also use graphics visualization and virtual reality tools to access the different resources via ESnet services. Members of the CMSN are located both at universities and the national laboratories and will require sustainable network services for using these resources.

3

Core Program Requirements

Biological Research

ESnet

Program Description

Within DOE's Office of Biological and Environmental Research (BER), biological research is funded and managed primarily through two divisions: Life Sciences and Medical Sciences *(http://www.er.doe.gov/production/ober/ober_top.html)*. Additional programs are supported by the Environmental Sciences Division of BER, including studies of microbial systems, the Natural and Accelerated Bioremediation Research (NABIR) program, and the Climate Change Technology Initiative (CCTI).

These BER programs advance fundamental biological science at the intersection of biology with the physical, computational, and engineering sciences. They involve several hundred scientists and highly skilled technicians, as well many more students and associated faculty members. These researchers work in an interdisciplinary environment, operating a number of special facilities for the biomedical research and development community and collaborating with laboratories worldwide. The DOE biological research program is an integral part of the global biomedical science and technology enterprise. Progress in our understanding of biological systems depends critically on information technology, both in computational sciences and networking infrastructures. The biological research program supported by BER leverages special capabilities at the national laboratories in the following focus areas.

Functional genomics. Producing genome sequences is just one of the first steps toward a more complete understanding of biological systems. The BER program is now looking at genome-scale approaches to determine the function of genes and systems of genes and to decipher other important signals encoded in DNA sequences. This growing thrust includes comparative genomics, which is based on comparisons of mouse and human sequences and analysis of other appropriate model organisms.

Microbial cells. BER has been a leader in sequencing microbial genomes and is pursuing a major new thrust in which the continuing work in microbial sequencing also supports efforts understand the complex mechanisms involved in a living microbial cell *(http://www.sc.doe.gov/production/ober/microbial.html)*.

Low-dose research. BER has pioneered research to understand the health effects associated with exposure to very low levels of ionizing radiation.

Experimental and computational structural biology. BER has unique facilities for experimental structural studies, among them synchrotron-based X-ray sources and neutron sources, world-renowned resources for biological mass spectrometry, and high-performance nuclear magnetic resonance (NMR) equipment *(http://www.sc.doe.gov/production/ober/msd_struct_bio.html)*. DOE also has a large computing infrastructure and capabilities for large-scale automation. These experimental and computational resources are now being integrated with gene sequencing and functional genomics studies involving human, mouse, and microbial systems. The result is a powerful new approach to biology that requires a state-of-the-art computer networking infrastructure.

Bioinformatics. BER has been a leader in developing analysis tools for genome sequence data, repositories for genome sequence information *(http://compbio.ornl.gov/channel/index.html)*, and genome information systems that serve as a resource for the research community and the general public *(http://www.ornl.gov/hgmis/)*. Such bioinformatics resources will be a critically important component of the BER biological research program, especially as complete sequences become available for more and more genomes and as the quest to understand biological systems reaches new levels of complexity.

Biomedical engineering. BER is developing a biomedical engineering program that is expected to grow in FY 2001 and beyond.

Current Program Requirements

Many of the research areas described above have significant networking requirements and could not exist in their present form without the Internet. The biology community relies on the network infrastructure to conduct collaborative research, for rapid access to many information resources as an integral part of the research effort, and for effective communications with the public. Two examples will serve to illustrate how BER programs use network computing. The first is the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL), a major user facility supported by BER *(http://www.emsl.pnl.gov:2080/)*. It includes a collaboratory providing remote access to state-of-the-art analytical instrumentation that requires appropriate bandwidth and connectivity provided by ESnet *(http://www.emsl.pnl.gov:2080/docs/collab/)*.

This microbe, Rhodopseudomonas palustris, *can degrade and recycle components of the woody tissues of plants. Because of its intimate involvement in carbon management and recycling (key issues in the question of global warming), this bacterium was selected by the DOE Carbon Management Program for genome sequencing.*

The second is the Human Genome Project and other large-scale genome sequencing programs. These global efforts are producing data sets that are freely accessible. Sequence data contain information about all the genes in a specific organism (the "parts list"). Such data also carry a rich set of signals that control expression of single genes and systems of genes in space (i.e., tissue) and time during the life of the organism (the "manual" on how the parts fit together and are to be used). Unless this information is derived from the sequence and added to it in a way that is useful for biological researchers, the value of sequence data is relatively insignificant. Some analysis tools are quite well developed (e.g., gene finding) and represent mature technology. Others continue to present significant research challenges, such as tools for analysis of regulatory regions and for comparisons across genomes. Many such tools are being used routinely via the World Wide Web. Often these applications draw on terascale computing resources linked via ESnet, and data traffic is likely to increase for the foreseeable future.

Traffic is already growing rapidly. The DOE Genome Channel resource *(http://compbio.ornl.gov/)* had approximately 4 million data transfers in May 2000, up from 1 million in March 2000. The DOE Human Genome Project Information (HGPI) Web site *(http://www.ornl.gov/hgmis/)* statistics show the following usage levels (as of 2000): 3.95 million text pages viewed per year, 1 million user sessions per year, an average session length of 12 minutes, and more than 8,000 links into the site from external pages. The approximate number of text file hits per year at this site has grown from 23,000 in 1995 to a projected 7 million in 2000, having approximately doubled every year since 1996.

DOE genome sequencing facility supports Human Genome Project

DOE's Office of Biological and Environmental Research initiated the Human Genome Project and has been a partner in the global sequencing effort. Through its Joint Genome Institute, DOE has established the Production Sequencing Facility for genome sequencing (http://www.jgi.doe.gov/). This facility can produce the draft sequence of a 500-million-base genome (e.g., a small vertebrate genome or large mammalian chromosome) in about 10 months (from eight sequencing passes). This facility recently met a major milestone for the DOE human genome program by completing the draft sequence of human chromosomes 5, 16, and 19 (see announcement, http://www.ornl.gov/hgmis/project/51619jgi.html). It is also performing high-throughput sequencing of microbial genomes.

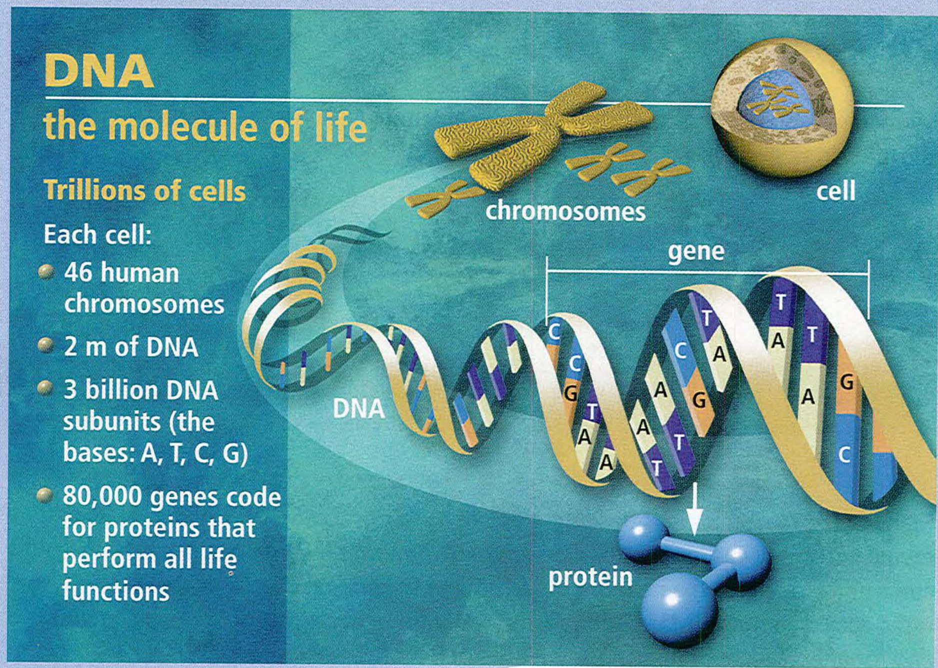

DOE press release, April 13, 2000

DOE Joint Genome Institute Announces Draft Sequence of Human Chromosomes 5, 16, and 19

U.S. Secretary of Energy Bill Richardson announced today that researchers at the Department of Energy's Joint Genome Institute in Walnut Creek, California, have decoded in draft form the genetic information on human chromosomes 5, 16 and 19. The chromosomes contain an estimated 10-15,000 genes, including those whose defects may lead to genetically linked diseases such as certain forms of kidney disease, prostate and colorectal cancer, leukemia, hypertension, diabetes and atherosclerosis.

"Three chapters in the reference book of human life are nearly complete," said Secretary Richardson. "Scientists can already mine this treasure trove of information for the advances it may bring in our basic understanding of life as well as applications such as diagnosing, treating and eventually preventing disease." Richardson made his remarks at the 25th Annual American Association for the Advancement of Science Colloquium on Science and Technology Policy in Washington, D.C.

Energy researchers contain more than 300 million base pairs, or an estimated 11 percent of the total human genome. To date the researchers have sequenced a working draft of the three chromosomes, leaving some scattered gaps in less gene-rich areas. Institute researchers will continue to improve both the completeness and accuracy of the genetic information as they produce the final sequence of the chromosomes over the next several years.

Chromosome 5 contains an estimated 194 million bases, or about 6 percent of the human genome. Disease-linked genes on this chromosome include those for colorectal cancer, basal cell carcinoma, acute myelogenous leukemia, salt-resistant hypertension and a type of dwarfism. Chromosome 16 contains about 98 million bases, or about 3 percent of the human genome. Studies have implicated genes on this chromosome in the development of breast and some prostate cancer, Crohn's disease and adult polycystic kidney disease, which affects an estimated five million people worldwide. Half the affected people require dialysis or kidney transplant. Chromosome 19 contains 60 million bases, or about 2 percent of the human

in Berkeley Calif., the Lawrence Livermore National Laboratory in Livermore, Calif. and the Los Alamos National Laboratory in Los Alamos, New Mexico. The University of California manages all three laboratories for the Energy Department. The department's Oak Ridge National Laboratory is analyzing the sequences computationally to identify the locations of genes within the sequences.

The Department of Energy began the Human Genome Project in 1986 to map and determine the complete DNA sequence of the human genome. The project's ultimate goal is to discover the 80 to 100,000 human genes on the 23 pairs of chromosomes and enable biologists to study them in detail. The publicly funded effort, now an international research project, divides responsibility for sequencing the chromosomes among the participating sequencing centers. Other participants include the U.S. National Institutes of Health and the Wellcome Trust in England which jointly announced in December the completion of the first final sequence for a chromosome, chromosome 22. The international project expects to complete the draft for the entire human genome in June, with a final sequence available on or before 2003. Last month, the international consortium completed the draft

Genome research depends on both data-intensive and computation-intensive analyses

The ability to study the content and structure of genomes and the relationship between the structure and function of proteins makes it possible to study life in its most fundamental details, from molecules to cells to organisms. High-performance networked computing is central to this work. The results of this research will profoundly affect our understanding of human health, our approach to health care, and our understanding of the relationships between human health and the environment.

Gene function. *The relationships between genetic material (genotype) and its physical expression (phenotype) in an organism are highly complex and, for the most part, poorly understood. To determine how a gene functions, computational analyses must be carefully combined with experimental approaches. Model organisms, such as the mouse, microbes, and other organisms, are being used to study gene function in the whole organism. Researchers work with these models by knocking out genes, inserting genes, or systematically causing mutations and then observing the resulting phenotypes.*

These efforts produce multi-dimensional and diverse sets of data, including, for example, complete three-dimensional imaging data for whole animals or organs and two-dimensional gene expression data as functions of time, tissue of origin, and experimental treatment. Integrating the data, deriving useful information for biomedical researchers, and making this information available to the research community represent significant research and computing challenges. Such tasks are being addressed by researchers in the BER program and by other groups in the global biology research enterprise. Access to the information resources produced by these efforts involves significant use of ESnet.

Protein function. *The genome encodes specifications for synthesis of proteins, the basic building blocks of the machinery of living organisms. 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 important to understanding how it performs its function.*

Currently, structural information can be gained at various levels of resolution through experiments involving X-ray or neutron sources and by NMR spectroscopy. Mass spectrometry represents another important analytical tool to characterize proteins. Computational methods for representing and predicting protein structure are improving in their capabilities and can now provide useful information in a growing number of cases. By combining experimental data with computational algorithms, researchers have created a powerful approach to determining structure. Beyond 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 such studies also depend upon a combination of experimental and computational methods.

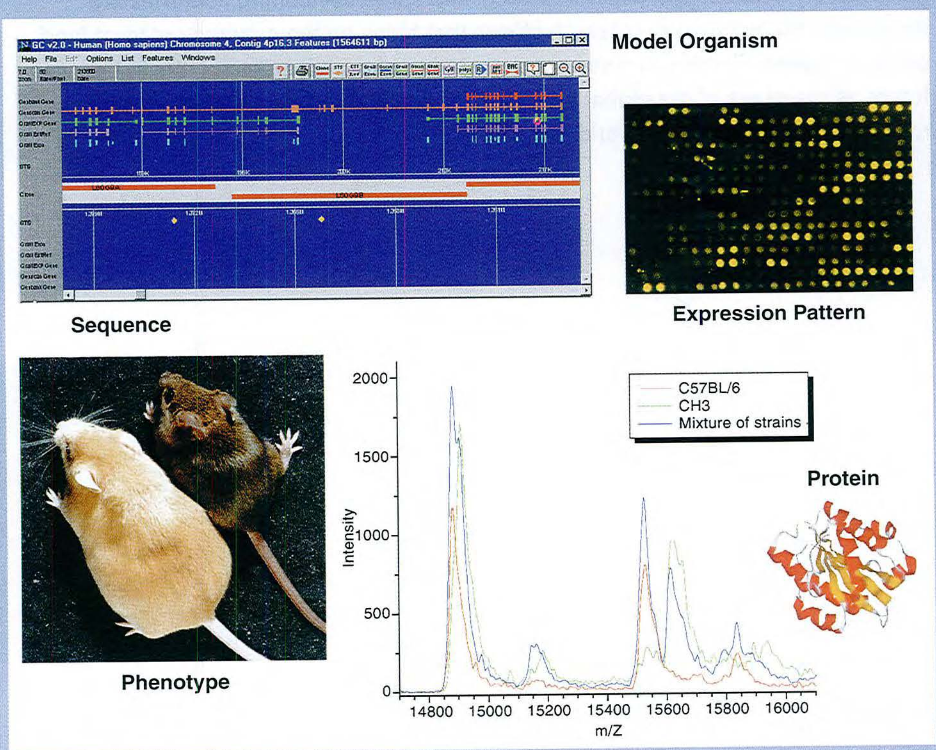

Computing resources are used to identify genome sequences, gene expression patterns, and protein structures for various model organisms, such as the mouse. The results help researchers understand the relationship between the genotype (genetic material) and phenotype (external characteristics of an individual).

Future Program Requirements

Biological research increasingly involves large facilities and resources that are made available to large numbers of remote users. The success of this enterprise depends on appropriate networking infrastructure and services that provide integrated views of the many data sets and information resources being created and updated as biological research advances rapidly.

The network infrastructure is an integral part of the biological research environment. 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 needs of the research programs involved. Numerous, high-capacity connections among major database locations and research centers are essential to allow researchers to access individual pieces of information and to allow database managers to cross-reference and use information in other databases.

Access to databases in foreign countries is essential to U.S. researchers. Current research takes place principally in the industrialized countries, and biology information resources are being maintained in these countries. However, this is one of the most rapidly expanding areas of research, and some developing countries are developing biotechnology research efforts that focus on the genomes of local food crops and pests. These countries are expected to play an increasing role in the near future as members of the global biological research enterprise, placing further requirements on global networks.

3 Core Program Requirements

Environmental Research

ESnet

Program Description

The Environmental Sciences Division (ESD) within the Office of Biological and Environmental Research supports basic research in two general program areas: Environmental Processes–Global Change Research and Environmental Remediation Research *(http://www.sc.doe.gov/production/ober/esdabout.html)*. A dependable computing network that provides adequate bandwidth, speed, and service is fundamental to research in ESD.

Environmental Processes–Global Change Research

Research and modeling in the Environmental Processes–Global Change Research area help scientists (1) improve their understanding of factors that affect the earth's radiant-energy balance; (2) accurately predict climate changes that could result from increased atmospheric concentrations of greenhouse gases; (3) quantify sources and sinks of energy-related greenhouse gases, especially carbon dioxide; and (4) improve the scientific basis for assessing the potential consequences of climate changes.

The research focuses on understanding the basic chemical, physical, and biological processes that take place in the earth's atmosphere, land, and oceans and how these processes can be affected by energy production and use. A major part of the research is designed to provide data that will enable researchers to objectively assess the potential for global warming and its possible consequences. This comprehensive program area emphasizes (1) the study of the radiation balance from the surface of the earth to the top of the atmosphere (including the role of clouds) and (2) enhancements to quantitative models so they can be used to predict climate change at global and regional levels.

Specific global change research programs, such as the Atmospheric Radiation Measurement (ARM) Program (see page 87) and Climate Change Prediction Program (CCPP), use ESnet heavily in day-to-day data collection, processing, and analyses. These programs depend on network technology to facilitate collaborations among scientists as they take measurements, analyze results, make observations, and run models.

ESnet is also valuable to researchers working in other ESD atmospheric science programs. To obtain a comprehensive understanding of the atmospheric processes that control the transport, transformation, and fate of energy-related chemicals and particulate matter, these researchers require access to an infrastructure of distributed computing and information processing capabilities. They can use ESnet to conduct high-quality analyses of complex data and to integrate and synthesize the information they need to evaluate environmental issues. A wide range of scientists and other participants from DOE laboratories, other federal laboratories, colleges, universities, and private industry are involved in such efforts.

Environmental Remediation Research

Programs in the Environmental Remediation Research area focus on understanding fundamental physical, chemical, geological, and biological processes. Results are marshaled to develop and advance new, effective, and efficient processes for remediating and restoring the nation's nuclear weapons production sites. Two programs in this area, the Microbial Genome Program and the Natural and Accelerated Bioremediation Research Program, use ESnet heavily to share large amounts

Atmospheric Radiation Measurement (ARM) Program

The ARM Program is a key component of DOE's research strategy for addressing concerns about global climate change and the increase in greenhouse gas emissions. The program is an observational and analytical research effort that focuses on collecting data needed to improve general circulation models (GCMs) and provide reliable simulations of regional and long-term climate change (http://www.arm.gov/).

Instruments at the ARM site on Nauru monitor clouds over the western Pacific Ocean

Its present objectives are to acquire data that will help researchers understand the role of clouds in climate and to ensure that this knowledge is reflected in the appropriate GCMs. Collaborations with researchers in climate change programs at other agencies are encouraged, and the ARM Program participates in the United States Global Change Research Program.

The ARM Program both depends upon and poses challenges for ESnet for reasons of geography and data volume. The program currently operates three sites: on the Southern Great Plains (Oklahoma), in the tropical western Pacific Ocean, and the North Slope of Alaska. Each site has multiple measurement locations. The first measurement location was established in 1992 and the most recent in 1999.

Among the data-intensive aspects of the ARM Program is the Unmanned Aerospace Vehicle (UAV) Program. This program provides data from instruments in planes flown for long periods at the top of the tropopause, roughly 6-8 miles (10-13 kilometers) above an ARM site. Measurements made by these planes and other instruments include vertical profiles of temperature, water vapor, trace gases, aerosols, and solar and infrared radiation.

Data from each site are being integrated into a central data management facility at Pacific Northwest National Laboratory. The data are available to the ARM Science Team in near real time and tailored to each Science Team project. The data are also available to the general science community through a data archive established at Oak Ridge National Laboratory. The 800 registered users of ARM archive data are widely dispersed, spanning 25 countries.

The Altus UAV serves as an airborne measurement platform of biological information, computational models, and integration tools. Moreover, the Environmental Molecular Sciences Laboratory (EMSL) depends highly on robust networking technology to support the collaborative nature of its research and its use of automated instrumentation in problem-solving environments. EMSL is the only national collaborative user facility that focuses on DOE's environmental mission, and it is a key element in supporting the DOE Office of Biological and Environmental Research's commitment to environmental remediation.

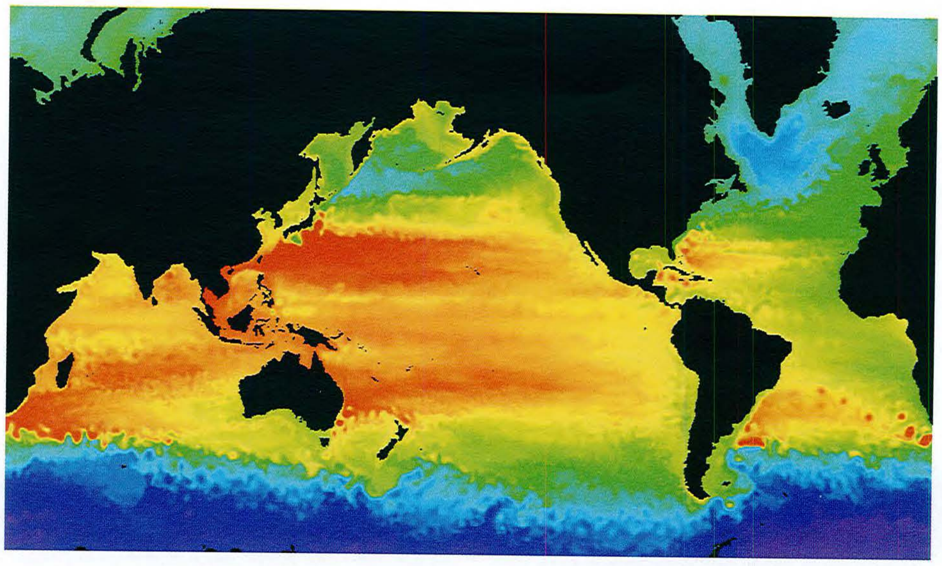

Highly detailed global climate modeling requires network access to remote supercomputers

National User Facility: William R. Wiley Environmental Molecular Sciences Laboratory (EMSL)

EMSL is a new national collaborative user facility located at Pacific Northwest National Laboratory (http://www.emsl.pnl.gov:2080/). It provides collocated facilities for multidisciplinary research in molecular-level environmental sciences.

This world-class scientific center provides strong interdisciplinary research environments, with links to users in the government, industry, and academia.

Many computational and scientific instrument systems in EMSL focus on helping researchers develop a molecular-level understanding of the physical, chemical, and biological processes that underlie (1) remediation of contaminated soils and groundwater, (2) processing and disposal of stored wastes, and (3) the effects on human health and the ecology from exposure to pollutants. In such work, geographically distributed and interdisciplinary teams of investigators collaborate in problem-solving sessions. EMSL relies heavily on networking technology to support its research programs.

EMSL employs a "laboratory without walls" concept to provide access to its equipment and data

Current Program Requirements

The flow, collection, processing, and delivery of ARM data are made possible by ESnet. Furthermore, ESnet forms the communication foundation that enables ARM and similar programs to build distributed systems and teams that can focus on solving specific problems. As a result, state-of-the-art network connections must be maintained to the ARM data centers at ORNL and PNNL, as well as the data center used in the CCPP. Access to ESnet is also very important to user facilities, specifically EMSL.

ESnet uses satellite links to connect to some remote sites

The ARM Program and CCPP are currently served well by DS-3 (45 Mbps) or OC-3 (155 Mbps) connectivity. ARM uses dedicated T1 links for data collection and communication with the Southern Great Plains and North Slope of Alaska. To communicate with the tropical Western Pacific site, an Inmarsat-B 64-kbps satellite link or Geostationary Operational Environmental Satellite (GOES) channel is used.

In 2001, the amount of data flowing through the ARM computing environment will grow to approximately 25 gigabytes per day. And this number does not include the quantities of data to be delivered to the ARM Science Team or to the general scientific community, both of which vary depending on research needs. Currently, the ARM Program's data holdings exceed 10 terabytes and are expected to grow to more than 20 terabytes by the end of 2003. The program depends on ESnet to move that data.

As indicated in the following diagram, the ARM computing infrastructure is geographically dispersed. Development, maintenance, and operation of the instrument and computing environment is the responsibility of the Engineering and Operations Teams of ARM. These teams are geographically dispersed, being composed of staff from several DOE national laboratories and research sites.

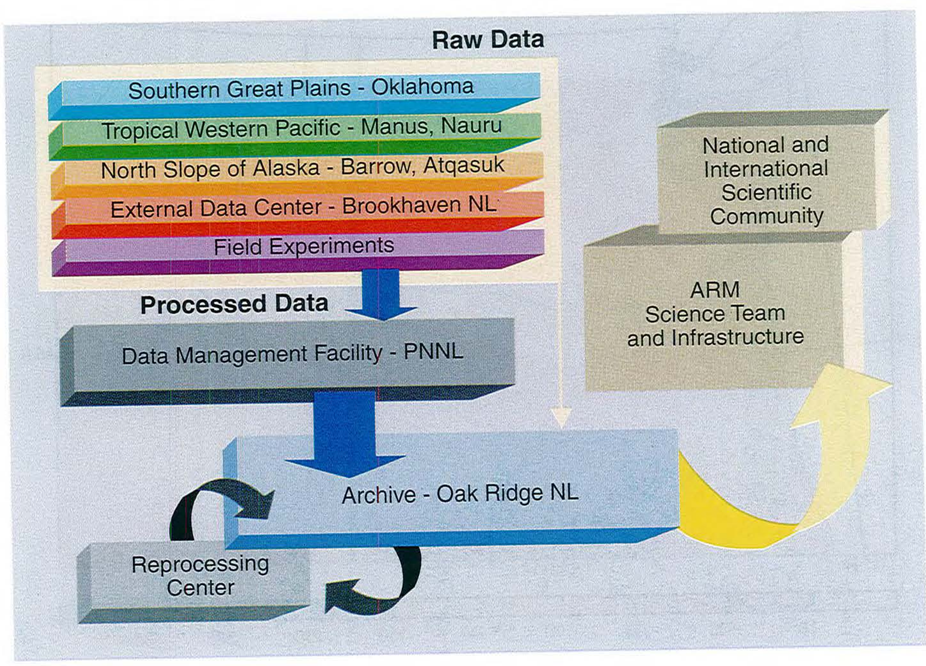

ARM data are collected, processed, and delivered to the user community in near real time

Future Program Requirements

Research using ARM data is expected to continue for several decades after the data collection program ends. These research studies will use large quantities of archival ARM data, which will be used for a variety of applications. For example, data will be needed to perform analyses designed to identify new and complex interactions between radiation energy and water (cloud) components. ARM data will also be needed for setting model input, comparing climate models, and validating information from NASA Earth Observing System (EOS) satellites. For this work, researchers will need to store and process (e.g., compare, summarize) tens of terabytes of data from field observations, satellites and remote sensors, and model output. Since 1992, 10 terabytes of data has been collected for the ARM Program, and 30 to 40 terabytes more is expected to be collected over the next five to seven years. Researchers located at many DOE, NASA, and National Oceanic and Atmospheric Administration (NOAA) laboratories and universities will use the data. Distributing the data and preliminary research results will require large quantities of information to be selected, transferred, and merged via the Internet.

As computing systems enter the era of teraflop speeds (trillions of floating-point operations per second) and the need to run climate models at finer resolutions becomes pressing, the ability to generate high volumes of model output will stress the resources for archiving and comparing this output. A typical high-resolution ocean model running on present-day computers having peak speeds in the 100-gigaflop range can generate a dozen multi-gigabyte files in a few hours, at an average rate of about 2 megabytes per second. At this rate, computing a century of simulated time would take more than a month to complete and would produce

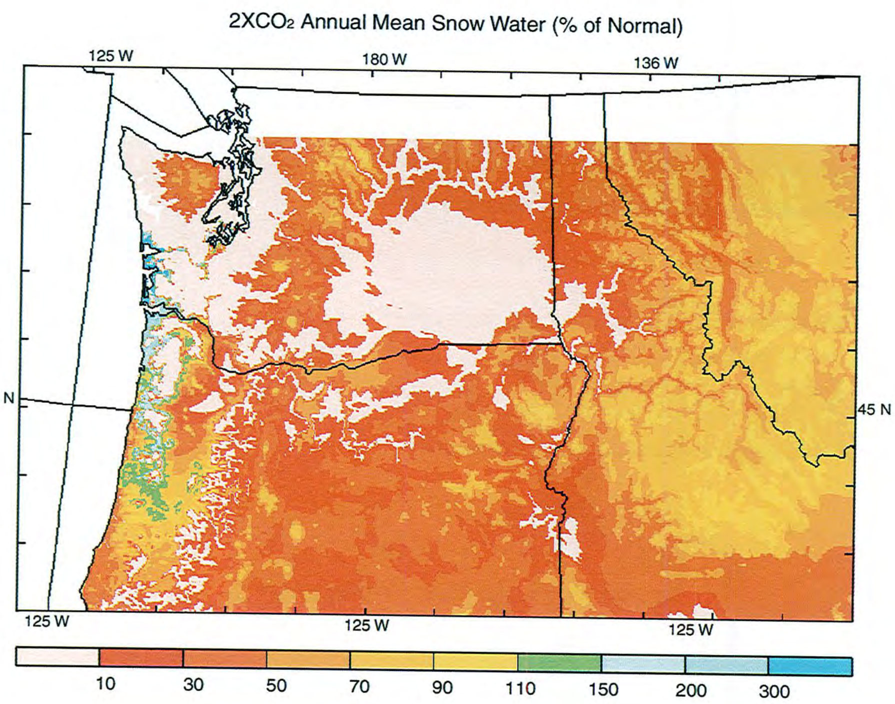

Image courtesy of Ruby Leung, PNNL

Regional model used to predict the impact of doubled CO_2 levels on climate in the Pacific Northwest. The figure shows the ratio of the annual mean snowpack under doubled CO_2 versus control conditions.

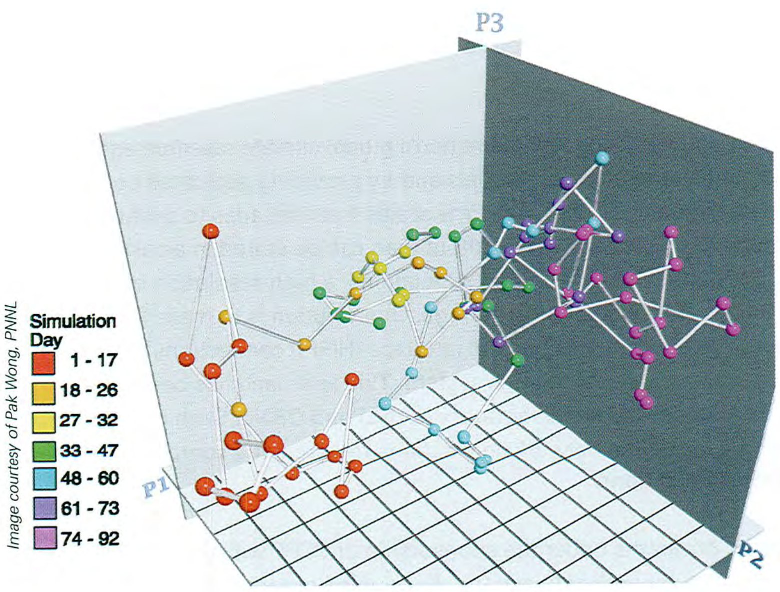

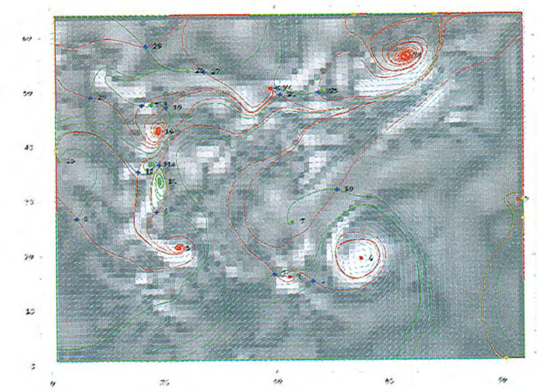

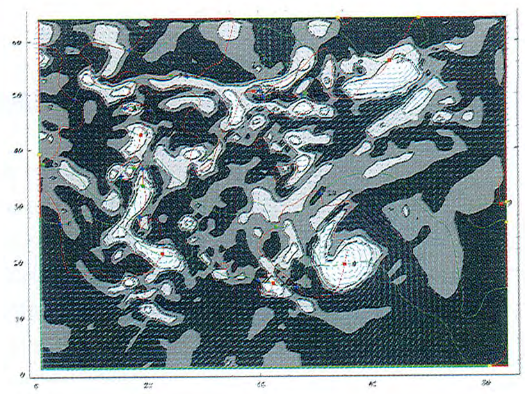

Data signatures and vector field simplification techniques: examples of new tools being evaluated for use in climate research

about 10 terabytes of output to be archived. Archival systems capable of storing hundreds of terabytes would be required to support calculations of this scale on a regular basis. Moving from a 100-gigaflop to a 1-teraflop system could multiply each of these figures tenfold, making petabyte archives essential. It is probable that machines capable of 20 to 40 teraflops will exist in 2003.

With regard to other atmospheric science programs, problem-solving environments for atmospheric, cloud, regional, and general climate models will require a distributed computing environment to

- Monitor real-time transfer of model output as it is being generated.
- Transfer model output after completion.
- Transfer processed, analyzed, or reduced data sets.
- Generate, transfer, or visualize user-selected reduced data sets.

The management of information, measurements, and measurement metadata over ESnet will be a key to the long-term collaborative success of atmospheric science programs. Tools for reviewing and analyzing metadata will be needed to provide a seamless way of integrating distributed data and metadata sources. The need will be scientifically driven, since knowledge of the quality of measurements and of collection conditions is fundamental to understanding and using the data.

Good navigation of metadata will be difficult not only because the generation and storage of the data and metadata holdings will be distributed but also because the information holdings will be so diverse. The display of user output resulting from metadata navigation and analysis will require the simultaneous processing of many complex requests for data and metadata that will span multiple measurement streams, geographic locations, and databases. The baseline performance of software tools is highly influenced by the performance of the Internet connections during transactions. Good Internet performance in terms of speed, dependability, and predictability is important, and in the next three years, it is predicted that OC-12 (622 Mbps) connectivity with differentiated services and intelligent middleware will be required.

DOE and the CCPP in particular support the nation's global climate research agenda by sponsoring the development of new models and by providing dedicated computer resources for climate simulations that span time scales from decades to centuries. However, these models can 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. One solution is to make local-area data transfers over high-performance parallel interface (HIPPI) connections between a parallel processor and an archival storage system. This approach has been implemented at both Los Alamos National Laboratory (LANL) and ORNL. Such configurations are typical of the local-area network connections required at a computer center that supports climate simulations.

The climate modeling programs under the sponsorship of CCPP will also require state-of-the-art wide-area network connections. As coupled oceanic and atmospheric climate models continue to be developed, enhanced connections between the national laboratories will become important, in particular, the connection between the high-performance computing research centers at Lawrence Livermore National Laboratory (LLNL) and ORNL. A T3 link would support 30 to 50 researchers who would be using the current generation of high-performance computers at these research centers. The examination of climate model output and the transfer of critical files among researchers require connections with higher speeds than can be achieved over T1 links. In fact, the remote display of graphic images could push even 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 associated with the National Center for Atmospheric Research in Boulder, Colorado; the Geophysical Fluid Dynamics Laboratory at Princeton, New Jersey; and the Climate Model Intercomparison Project at LLNL in California. In the future, ensuring the access of the climate modeling community in universities and national laboratories to these centers will continue to be important.

Available for sale to the public from:
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: 800-553-6847
Facsimile: 703-605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

Available electronically at http://www.doe.gov/bridge/

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: 865-576-8401
Facsimile: 865-576-5728
E-mail: reports@adonis.osti.gov

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

Publishing support provided by the Information and Publishing Division of Argonne National Laboratory

Printed on recycled paper

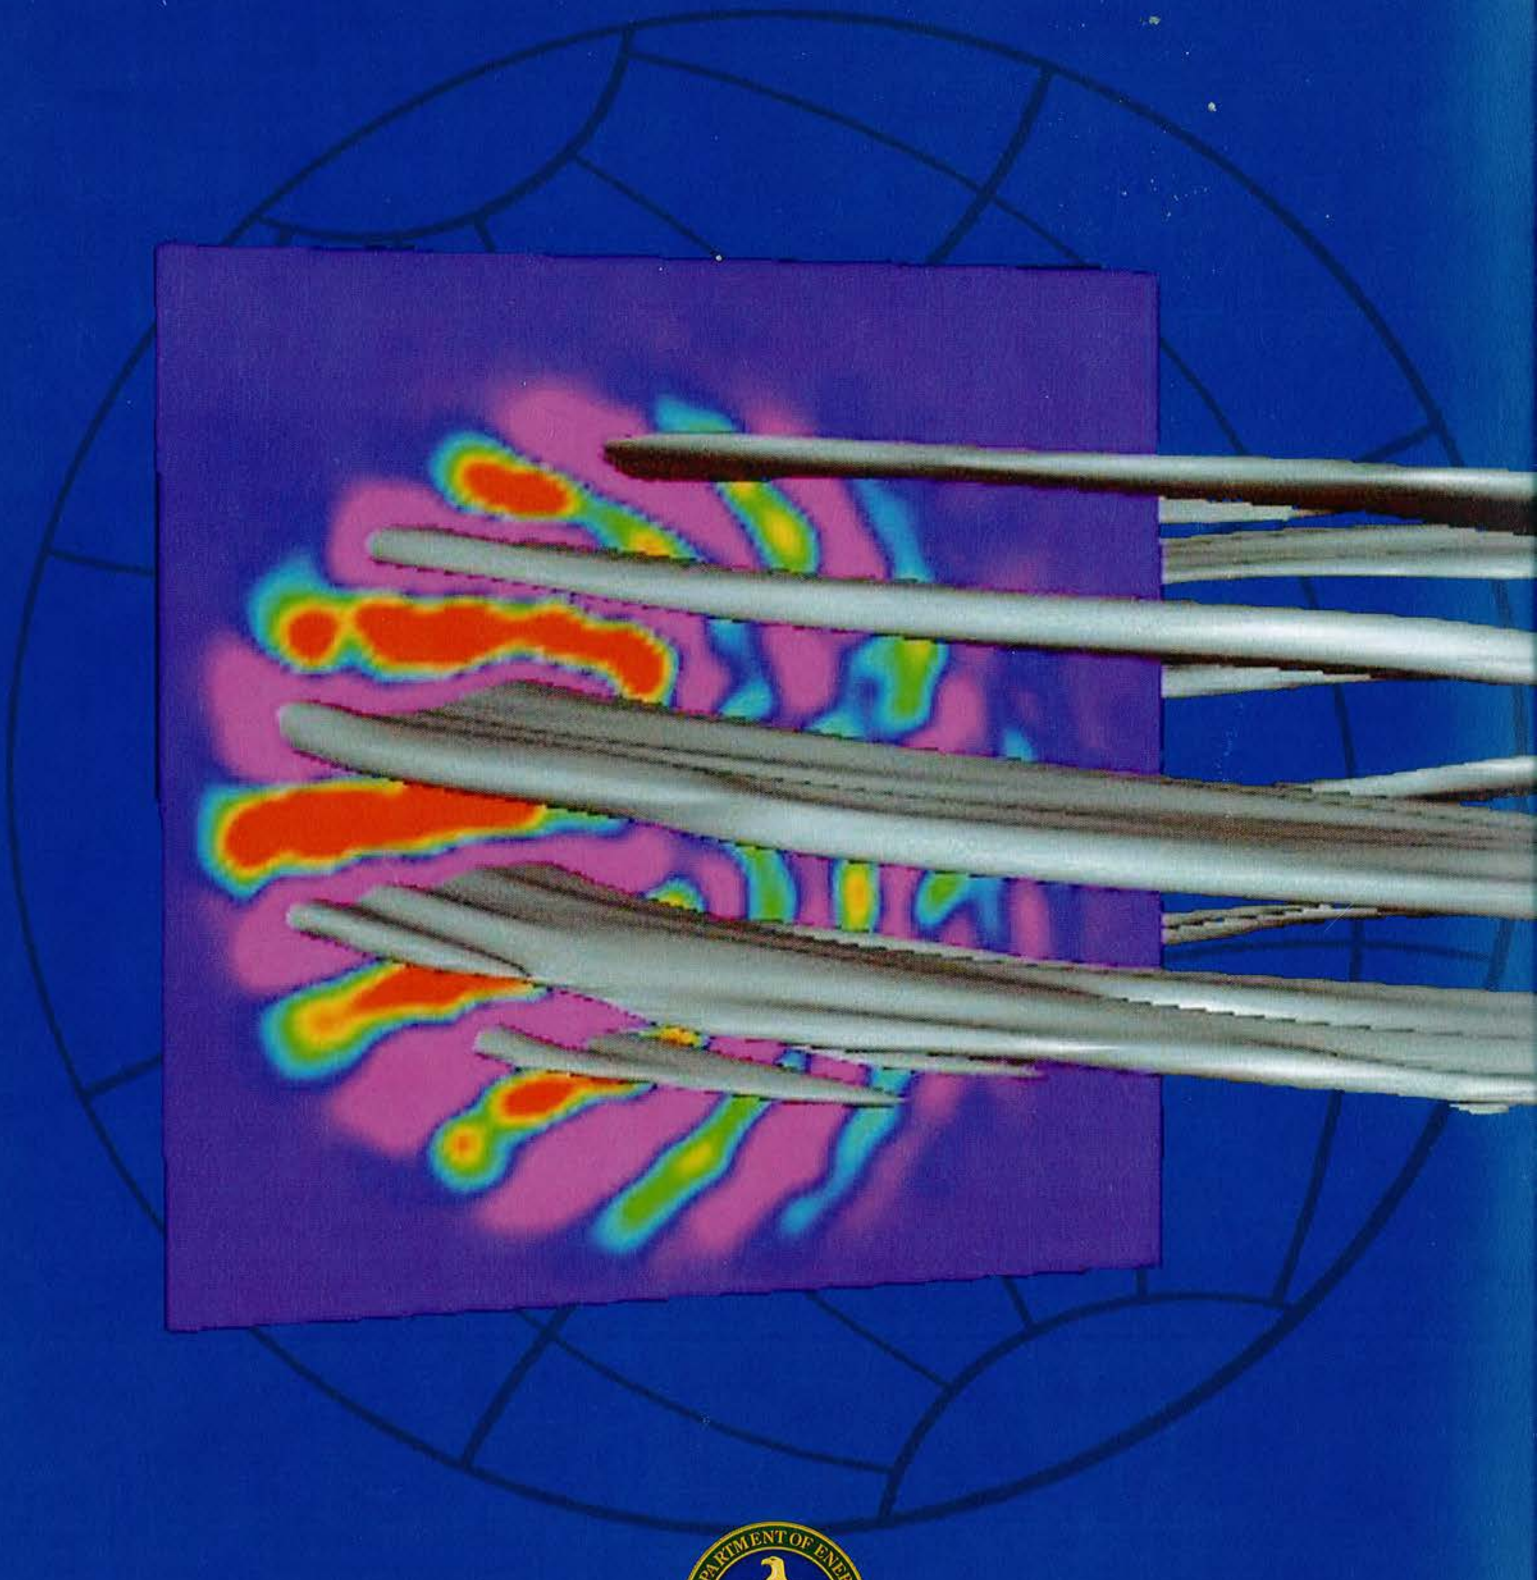

U.S. Department of Energy | Office of Science
Mathematical, Information, and Computational Sciences Division | Germantown, MD 20874