National Science Foundation
WHERE DISCOVERIES BEGIN

CYBERINFRASTRUCTURE:::::::::::::::: A SPECIAL REPORT

A GRAND CONVERGENCE
ESSENTIAL, NOT OPTIONAL
DISCOVERY, LEARNING AND LEADERSHIP
CLASSROOM RESOURCES

This animation depicts a hypothetical researcher using a well-established cyberinfrastructure to apply the power of the world's highest-performance digital resources

Credit: Animation by Nicolle Rager, National Science Foundation. Simulation Image by John Dubinski, Department of Astronomy and Astrophysics and Canadian Institute for Theoretical Astrophysics, University of Toronto. Telescope Image by Todd Boroson, NOAO/AURA/NSF

HPWREN's relay on Palomar Mountain, with smoke from the July 2003 Coyote wildfire rising in the distance. The

A Grand Convergence

Cyberinfrastructure is an idea that has emerged over the past decade or so from some basic technological realities:

- **The Internet** has become far more than an e-mail carrier;
- **Microchip technology** is filling our homes, our cars, and even our pockets with ever-shrinking computing devices; and
- **Databases** are piling up information faster than anyone can make sense of it all.

The convergence of these and other trends have led researchers to envision a tightly integrated, planet-wide grid of computing, information, networking and sensor resources--a grid that we could one day tap as easily as we now use a power socket or a water faucet.

That analogy is the source of the name, "cyber-infrastructure," and is quite apt. Modern researchers are coming to depend on cyberinfrastructure, even in its current embryonic form, in much the same way that modern society relies on its physical infrastructure to provide electricity, water, telephones and roads. In many disciplines, moreover, the results promise to be as revolutionary as the coming of water and electric power was for our cities.

With a well-established cyberinfrastructure, individual researchers will have the power of the world's highest-performance digital resources at their disposal. And teams of researchers will attempt to answer questions that had previously been unapproachable because the requirements were too much, too hard, too long or too complex.

"New instrumentation, data-handling and computation capabilities will expand the horizons of what scientists can study and understand," said Margaret Leinen, head of NSF's Geosciences directorate. "Cyberinfrastructure is empowering a new generation of researchers in their quest to unravel how the world around us works."

In environmental science for example, cyberinfrastructure combines computing, information management, networking and intelligent sensing systems into powerful tools for investigating the natural world and the human-built environment. The full complexity of the environment, from the molecular scale to the planetary, requires collecting and analyzing large

Cyberinfrastructure in Action

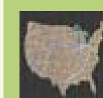

To iron out many of the technical wrinkles, the computer science community is collaborating with many potential beneficiaries in the scientific community and beyond to develop a number of ongoing NSF-supported cyberinfrastructure projects. Here are just a few:

***Extensible Terascale Facility*:** This NSF-supported collaboration, commonly called the TeraGrid, is a multi-year effort to prototype a widely shared, comprehensive cyberinfrastructure for academic research and education.

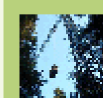

Networked Infomechanical Systems: In NIMS, robotic sensors suspended from a network of cables move themselves around to monitor, for example, a mountain stream ecosystem from the ground to the treetops.

Virus War Games: The DETER project allows researchers to study cybersecurity issues for Internet infrastructure and develop defense mechanisms in a controlled environment.

Data, Tools, Community: Unidata provides software infrastructure, data services and collaboration technologies as part of building a community for acquiring, organizing and using geoscience and geographic data.

Essential Privacy: Because of the proliferation of online sensitive data, the PORTIA project focuses on the technical challenges of handling sensitive data and the policy and legal issues facing data subjects, owners and users.

Crisis Response: Project RESCUE is transforming the way responding organizations gather, manage, use and

HPWREN team established a wireless network link to the California Department of Forestry and Fire Protection's operations camp for fighting the Coyote fire in northeastern San Diego County.

Credit: High Performance Wireless Research and Education Network (HPWREN)

volumes of data, performing experiments with computer models and bringing together collaborators from many disciplines. For example, as part of cyberinfrastructure, large-scale sensor networks or experimental instruments could feed real-time data streams into vast data grids where they can be funneled directly into ongoing computational simulations or shared as a community resource.

Nearly all scientific disciplines grapple with similar challenges in managing and analyzing vast amounts of data and complex processes across many scales. "The emerging cyberinfrastructure that NSF supports is a product of the scientific community's demands for and reliance on information and communications technologies," said Peter Freeman, head of NSF's Computer and Information Science and Engineering (CISE) directorate.

Cyberinfrastructure has become a common theme throughout NSF, and every directorate has funded or is exploring cyberinfrastructure-related projects. The NSF's larger goal for a national cyberinfrastructure is to provide the information technology and knowledge management resources needed to tackle the problems at the frontiers of all science and engineering disciplines.

disseminate information within emergency response networks and to the general public.

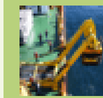

Wireless Without Limits: The HPWREN and ROADnet projects extend wireless research and education networks to remote astronomy observatories, Indian reservations and even ocean-going vessels.

By David Hart

A GRAND CONVERGENCE

ESSENTIAL, NOT OPTIONAL

DISCOVERY, LEARNING AND LEADERSHIP

CLASSROOM RESOURCES

Essential, Not Optional

"[E]nvironments and organizations, enabled by cyberinfrastructure, are increasingly required to address national and global priorities, such as understanding global climate change, protecting our natural environment, applying genomics-proteomics to human health, maintaining national security, mastering the world of nanotechnology, and predicting and protecting against natural and human disasters, as well as to address some of our most fundamental intellectual questions such as the formation of the universe and the fundamental character of matter."

That's how the NSF Blue Ribbon Advisory Committee in their report, Revolutionizing Science and Engineering Through Cyberinfrastructure, summarized the scientific need for cyberinfrastructure. The new capabilities are "essential, not optional, to the aspirations of research communities," the report states.

In the very near future, cyberinfrastructure will permit better forecasts of when and where earthquakes are likely to occur and how the ground will shake as a result. Robotic sensors will monitor ecosystem health or track pollutants in urban watersheds. Astronomers will mine for data as often as they peer into the heavens.

For scientists and engineers, according to the report, cyberinfrastructure has the potential to "revolutionize what they can do, how they do it, and who participates." The seeds of this revolution are seen in community-driven efforts, supported by NSF and other agencies, such as the Network for Earthquake Engineering Simulations (NEES), the Grid Physics Network (GriPhyN) and the National Virtual Observatory (NVO).

"Extensive grassroots activity by the scientific and engineering research community is creating and using cyberinfrastructure to empower the next wave of discovery," said Dan Atkins of the University of Michigan, chair of the advisory committee. "We're at a threshold where technology allows people, information, computational tools and research instruments to be connected on a global scale."

The report emphasizes the importance of acting quickly and the risks of failing to do so. The risks include lack of coordination, which could leave key data in irreconcilable formats; long-term failures to archive and curate data collected at great expense; and artificial barriers between disciplines

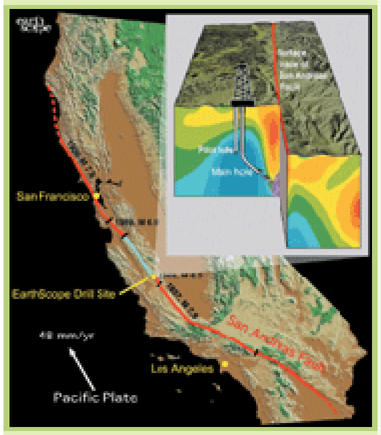

Deep within an active fault zone, EarthScope's San Andreas Fault Observatory at Depth will measure changes in rock properties before, during and after earthquakes. Linked to other EarthScope measurements at the surface, these direct observations will for the first time, monitor how an active fault and its environment respond to regional and local changes in stress. Recorded over a decade, this combination of measurements will provide important new insights on earthquake nucleation and rupture.

Credit: EarthScope

Grassroots Cyberinfrastructure

The demand for cyberinfrastructure has come from research communities that recognize the many ways it will allow them to push the scientific envelope. Here are a few of the disciplines and NSF-supported projects in the grassroots drive for cyberinfrastructure.

Physics: The Grid Physics Network (GriPhyN) and the international Virtual Data Grid Laboratory (iVDGL) are advancing data-intensive science for physics, cosmology and astrophysics.

Geoscience: EarthScope is investigating the evolution of the North American continent and the processescontrolling earthquakes and volcanic eruptions. The Geosciences Network (GEON) is integrating data, computation and software to weave together Earth sciences data and disciplines.

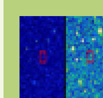

Astronomy: The National Virtual Observatory (NVO) is creating standards and high-performance computing tools for astronomical data collections that will make data easier to use, find and join with other data.

Engineering: The shake tables, reaction wall facilities, geotechnical centrifuges, tsunami wave tanks and mobile field equipment of the George E. Brown Jr. Network for Earthquake Engineering Simulation (NEES) are being integrated through the NEESgrid.

Meteorology: The Linked Environments for Atmospheric Discovery (LEAD) project will dramatically improve warnings of severe local weather such as thunderstorms and tornadoes by developing capabilities for on-demand detection, simulation and prediction.

Ecology: The Science Environment for Ecological built from incompatible tools and structures.

Advances in information technology and the mushrooming of cyberinfrastructure projects for specific fields demonstrate that NSF has a "once-in-a-generation opportunity," according to the committee, to lead the scientific and engineering community in the coordinated development and expansive use of cyberinfrastructure.

Knowledge (SEEK) project is a cyberinfrastructure for ecological, environmental and biodiversity research. The proposed National Ecological Observatory Network (NEON) will be a continental-scale research instrument for ecological studies.

Phylogeny: The Cyberinfrastructure for Phylogenetic Research (CIPRes) project is developing the technologies to reconstruct the tree of life starting from huge data sets containing hundreds of thousands of genomic sequences.

Discovery, Learning and Leadership

Cyberinfrastructure activities are the most recent developments in NSF's long history of leadership in providing advanced information technologies for the U.S. academic community.

NSF supported campus computing centers in the 1960s and established national supercomputer centers in the 1980s. In parallel, NSF established NSFnet in the mid-1980s, which evolved into today's commercial Internet; and in the 1990s, helped connect hundreds of institutions to advanced research networks. Also in the late 1990s, NSF established two Partnerships for Advanced Computational Infrastructure (PACI).

"Today's cyberinfrastructure exists thanks in part to the successful efforts of the PACI program," said Peter Freeman, head of NSF's Computer and Information Science and Engineering (CISE) directorate. "Through many activities, the two partnerships have nurtured and supported the growing demand by the scientific community for cyberinfrastructure."

PACI participants established the Extensible Terascale Facility (ETF), commonly called the TeraGrid, and planted the seeds of many discipline-specific cyberinfrastructure projects, supported by NSF and other agencies. Other NSF programs, including the Digital Libraries Initiative, the NSF Middleware Initiative and the Information Technology Research priority area also played key roles. NSF also has a history of collaboration with other federal agencies -- including NASA, the Department of Energy and others -- on cyberinfrastructure activities ranging from the Digital Libraries initiative to ETF to the Earth Systems Modeling Framework.

NSF's fiscal year 2005 budget request calls for $400 million for cyberinfrastructure-related activities across the foundation. In its recent reorganization, the CISE directorate created the Division of Shared Cyberinfrastructure to focus on such capabilities for a wide spectrum of science and engineering activities.

The National Center for Supercomputing Applications and the San Diego Supercomputer Center, the lead centers for the PACI partnerships, will continue to provide high-end computing resources and related services to the national community, at least through 2007. The ETF—on track to be commissioned in October 2004—will be widely shared for research on subjects ranging from galaxy formation to pollution cleanup

Three shake tables at the University of Nevada, Reno, were used to replicate the 1940 El Centro, Calif., "Imperial Valley" earthquake, which measured 7.1 on the Richter scale. A large beam, representing a model bridge, is mounted across the tables, which are shaken by large piston arms. The beam wobbled in response to the initial temblor and continued to shake for the duration of the quake. The bridge structure suffered no visible damage, but data from sensors was captured and transmitted to the NEESgrid infrastructure.

Credit: David Gehrig, NCSA

A History of Leadership

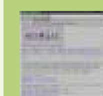

The following accounts provide more detail on NSF's history of leadership in computing and the programs that have contributed to existing and emerging cyberinfrastructure activities.

From Supercomputing to the TeraGrid: From the 1960s to today, NSF has supported the academic community with access to the highest-performance computing resources available.

A Brief History of NSF and the Internet: The Internet that we take for granted arose from a series of computer networking efforts funded by NSF and other agencies.

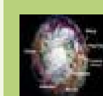

Networking for Tomorrow: NSF has long been a leader in cutting-edge network technologies and promotes activities to explore and engineer the Internet of the future.

Grand Challenges for Information Technology: As information technology and the Internet matured in the 1990s, NSF began the first of several initiatives to apply information technology to the most challenging research problems around.

Digital Libraries-Access to Human Knowledge: Digital libraries allow scientists, teachers, students and even software applications to access, explore, search and interact with vast data collections.

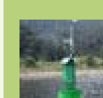

NMI in the Middle: Middleware is the "plug-and-play" interface for the many types of cyberinfrastructure resources, and the NSF Middleware Initiative is building a flexible and reliable middleware foundation.

through 2009.

NSF is continuing to solicit advice and input from the academic community on management of the emerging cyberinfrastructure and to inform NSF's development of future cyberinfrastructure-enhancing competitions. Internally, NSF has convened a cyberinfrastructure working group to explore challenges and opportunities in all science and engineering fields.

"NSF's goal is to ensure the development of a national cyberinfrastructure that is second to none," Freeman said. "The resulting widely-distributed, shared cyberinfrastructure will advance discovery, learning and innovation across the science and engineering enterprise."

More Special Reports

Mercury, the first TeraGrid cluster, deployed at the National Center for Supercomputing Applications (NCSA), runs on Intel's Itanium architecture. The 512-processor cluster has a peak performance of 2.7 teraflops (trillions of calculations per second).

Credit: NCSA

From Supercomputing to the TeraGrid

This history points out some of the landmarks along the route the National Science Foundation (NSF) has traveled to bring high-performance computing to the nation's researchers.

Credit: NCSA, University of Illinois at Urbana-Champaign

The Early History: 1960s-1980s. NSF's investment in the nation's computational infrastructure began modestly in the 1960s, when NSF funded a number of campus computing centers. That support was short-lived, however, and by the early 1980s, several reports from the scientific community noted a dramatic lack of advanced computing resources available to researchers at American universities. The most influential of these was a joint agency study edited by Peter Lax and released in December 1982.

The Lax Report led to the emergence of significant new NSF support for high-end computing, which in turn led directly to Supercomputer Centers.

The Supercomputer Centers: 1985-1997. NSF established five of these centers in 1985 and 1986:

- The Cornell Theory Center (CTC) at Cornell University,
- The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign,
- The Pittsburgh Supercomputing Center (PSC) at Carnegie Mellon University and the University of Pittsburgh,
- The San Diego Supercomputer Center (SDSC) at the University of California, San Diego,
- The John von Neumann Center at Princeton University.

For the next 12 years, these centers would continue to serve as cornerstones of the nation's high-performance computing and communications strategy. On the one hand, they helped push the limits of advanced computing hardware and software, even as they provided supercomputer access to a broad cross-section of academic researchers regardless of discipline or funding agency. And on the other, the centers were instrumental in advancing network infrastructure. In 1986, the centers and the NSF-supported National Center for Atmospheric Research (NCAR) in Colorado became the first nodes on the NSFNET backbone. From 1989 to 1995, the Illinois, Pittsburgh and San Diego centers helped push the frontiers of high-speed networking as participants in the then-bleeding-edge Gigabit Network Testbed Projects, which were supported by NSF and the Defense Advanced Research Projects Agency (DARPA). In 1995, after NSFNET was decommissioned, the centers became the first nodes on NSF's very-high-performance Backbone Network Service (vBNS) for research and education. (See "A Brief History of NSF and the Internet" and "Networking for Tomorrow" for more details.)

In 1990, following a review of the supercomputer centers program, NSF extended support for CTC, NCSA, PSC and SDSC through 1995. In 1994, that support was extended again for another two years, through 1997, while a task force chaired by Edward Hayes considered the future of the program.

Out of the recommendations of the Hayes Report came a new program designed to build on and replace the centers—

Partnerships for Advanced Computational Infrastructure: 1997-2004. The National Science Board announced the two PACI awardees in March 1997:

- The National Computational Science Alliance: a consortium led by NCSA, with

participation by Partners for Advanced Computational Services at Boston University, the University of Kentucky, the Ohio Supercomputer Center, the University of New Mexico, and the University of Wisconsin.
- The National Partnership for Advanced Computational Infrastructure (NPACI): a consortium led by SDSC, with participation by mid-range computing centers at Caltech, the University of Michigan, and the Texas Advanced Computing Center at the University of Texas at Austin.

In addition to leading-edge and mid-range sites, the PACI partnerships involved nearly 100 sites across the country in efforts to make more efficient use of high-end computing in all areas of science and engineering. The partnerships also collaborated on the Education, Outreach and Training (EOT) PACI.

The Alliance and NPACI continued to give academic researchers access to the most powerful computing resources available. These resources included the first academic teraflops system--a computer capable of 1 trillion operations per second--and some of the first large-scale Linux clusters for academia. At the same time, the partnerships were instrumental in fostering the maturation of grid computing and its widespread adoption by the scientific community and industry. Grid computing connects separate computing resources in order to apply their collective power to solve computationally intensive problems.

The PACI partners were involved in virtually every major grid-computing initiative, from the Grid Physics Network to the National Virtual Observatory to the George E. Brown, Jr. Network for Earthquake Engineering Simulation. The PACI partners were also driving forces in recognizing the critical scientific importance of and the technical challenges in accessing massive data collections. Following the sunset of the PACI program, NSF also continued core support for NCSA and SDSC to make more large-scale HPC resources available and to stimulate the expansion of cyberinfrastructure capabilities for the nation's scientists and engineers.

Terascale Initiatives: 2000-2004. In response to a 1999 report by the President's Information Technology Advisory Committee, NSF embarked on a series of "terascale" initiatives to acquire: (1) computers capable of trillions of operations per second (teraflops); (2) disk-based storage systems with capacities measured in trillions of bytes (terabytes); and (3) networks with bandwidths of billions of bits (gigabits) per second.

In 2000, the $36 million Terascale Computing System award to PSC supported the deployment of a computer (named LeMieux) capable of six trillion operations per second. When LeMieux went online in 2001, it was the most powerful U.S. system committed to general academic research. Five years later, it remains a highly productive system.

In 2001, NSF awarded $45 million to NCSA, SDSC, Argonne National Laboratory, and the Center for Advanced Computing Research (CACR) at California Institute of Technology, to establish a Distributed Terascale Facility (DTF). Aptly named the TeraGrid, this multi-year effort aimed to build and deploy the world's largest, fastest, most comprehensive, distributed infrastructure for general scientific research.

The initial TeraGrid specifications included computers capable of performing 11.6 teraflops, disk-storage systems with capacities of more than 450 terabytes of data, visualization systems, data collection, integrated via grid middleware and linked through a 40-gigabits-per-second optical network.

In 2002, NSF made a $35 million Extensible Terascale Facility (ETF) award to expand the initial TeraGrid to include PSC and integrate PSC's LeMieux system. Resources in the ETF give the national research community more than 20 teraflops of computing power distributed among the five sites and nearly one petabyte (one quadrillion bytes) of disk storage capacity.

To further expand the TeraGrid's capabilities, NSF made three Terascale Extensions awards totaling $10 million in 2003. The new awards funded high-speed networking connections to link the TeraGrid with resources at Indiana and Purdue Universities, Oak Ridge National Laboratory, and the Texas Advanced Computing Center at The University of Texas, Austin. Through these awards, the TeraGrid put neutron-scattering instruments, large data collections and other unique resources, as well as additional computing and visualization resources, within reach of the nation's research and education community.

In 2004, as a culmination of the DTF and ETF programs, the TeraGrid entered full production mode, providing coordinated, comprehensive services for general U.S. academic research.

The TeraGrid: 2005-2010. In August 2005, NSF's newly created Office of Cyberinfrastructure extended support for the TeraGrid with a $150 million set of awards for operation, user support and enhancement of the TeraGrid facility over the next five years. Using high-performance network connections, the TeraGrid now integrates high-performance computers, data resources and tools, and high-end experimental

Credit: Nicolle Rager-Fuller, National Science Foundation

facilities around the country. In early 2006, these integrated resources included more than 102 teraflops of computing capability and more than 15 petabytes (quadrillions of bytes) of online and archival data storage with rapid access and retrieval over high-performance networks. Through the TeraGrid, researchers can access more than 100 discipline-specific databases. With this combination of resources, the TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research.

Related Websites:
Lax Report (1982): http://www.pnl.gov/scales/archives.stm
Hayes Report (1995): http://www.nsf.gov/pubsys/ods/getpub.cfm?nsf9646
TeraGrid: http://teragrid.org/
NSF Office of Cyberinfrastructure: http://www.nsf.gov/dir/index.jsp?org=OCI
National Center for Atmospheric Research: http://www.ncar.ucar.edu/
Grid Physics Network: http://www.griphyn.org/
National Virtual Observatory: http://www.us-vo.org/

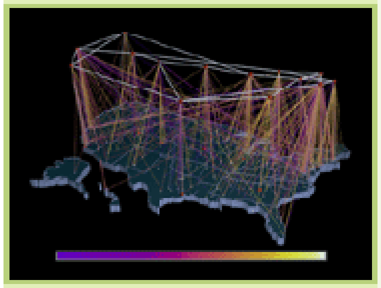

A visualization study of NSFNET backbone traffic in September 1991. The traffic volume ranges from zero bytes (purple) to 100 billion bytes (white). It represents data collected by Merit Network, Inc. In 1985, with the creation of the Supercomputer Centers program, NSF created NSFNET, a network that connected the five supercomputer centers and provided a network for research and education. Based on the ARPANET protocols, the NSFNET created a national backbone service.

Credit: NCSA

A Brief History of NSF and the Internet

The Internet that many of us take for granted today arose from a series of government-funded computer networking efforts, not the least of which was NSFNET. Along the way, NSF-supported researchers and students developed the first modern Web browser. These developments are just two of the significant milestones in NSF's contribution to the modern Internet.

Credit: NCSA, University of Illinois Board of Trustees

The Early Years: 1960s-1985. In 1969, the precursor to the Internet began with the U.S. Defense Department's ARPAnet. ARPA-funded researchers developed many of the protocols still used for most Internet communication. Several other agencies also developed networks so their researchers could communicate and share data. In 1981, for example, NSF provided a grant to establish the Computer Science Network (CSNET) for university computer scientists.

In 1985, NSF considered how it could provide greater access to the high-end computing resources at its recently established supercomputer centers. Because NSF intended the supercomputers to be shared by scientists and engineers around the country, any viable solution had to link many research universities to the centers.

NSFNET: 1986-1995. NSFNET went online in 1986 and connected the supercomputer centers at 56,000 bits per second—the speed of a typical computer modem today. In a short time, the network became congested and, by 1988, its links were upgraded to 1.5 megabits per second. A variety of regional research and education networks, supported in part by NSF, were connected to the NSFNET backbone, thus extending the Internet's reach throughout the United States.

Creation of NSFNET was an intellectual leap. It was the first large-scale implementation of Internet technologies in a complex environment of many independently operated networks. NSFNET forced the Internet community to iron out technical issues arising from the rapidly increasing number of computers and address many practical details of operations, management and conformance.

Throughout its existence, NSFNET carried, at no cost to institutions, any U.S. research and education traffic that could reach it. At the same time, the number of Internet-connected computers grew from 2,000 in 1985 to more than 2 million in 1993. To handle the increasing data traffic, the NSFNET backbone became the first national 45-megabits-per-second Internet network in 1991.

The history of NSFNET and NSF's supercomputing centers also overlapped with the rise of personal computers and the launch of the World Wide Web in 1991 by Tim Berners-Lee and colleagues at CERN, the European Organisation for Nuclear Research, in Geneva, Switzerland. The NSF centers developed many tools for organizing, locating and navigating through information, including one of the first widely used Web server applications.

But perhaps the most spectacular success was Mosaic, the first freely available Web browser to allow Web pages to include both graphics and text, which was developed in 1993 by students and staff working at the NSF-supported National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign. In less than 18 months, NCSA Mosaic became the Web "browser of choice" for more than a million users and set off an exponential growth in the number of Web servers as well as Web surfers. Mosaic was the progenitor of modern browsers such as Microsoft Internet Explorer and Netscape Navigator.

Internet Privatization: 1993-1998. Commercial firms noted the popularity and effectiveness of the growing Internet and built their own networks. The proliferation of private suppliers led to an NSF solicitation in 1993 that outlined a new Internet architecture that largely remains in place today.

From that solicitation, NSF awarded contracts in 1995 for three network access points, to provide connections between commercial networks, and one routing arbiter, to ensure an orderly exchange of traffic across the Internet. In addition, NSF signed a cooperative agreement to establish the next-generation very-high-performance Backbone Network Service. A more prominent milestone was the decommissioning of the NSFNET backbone in April 1995.

In the years following NSFNET, NSF helped navigate the road to a self-governing and commercially viable Internet during a period of remarkable growth. The most visible, and most contentious, component of the Internet transition was the registration of domain names. Domain name registration associates a human-readable character string (such as "nsf.gov") with Internet Protocol (IP) addresses, which computers use to locate one another.

The Department of Defense funded early registration efforts because most registrants were military users and awardees. In the early 1990s, the Federal Networking Council (a group of government agencies involved in networking) asked NSF to assume responsibility for non-military Internet registration. When NSF awarded a five-year agreement for this service to Network Solutions, Inc. (NSI), in 1993, there were 7,500 domain names.

In September 1995, as the demand for Internet registration became largely commercial (97 percent) and grew by orders of magnitude, the NSF authorized NSI to charge a fee for domain name registration. Previously, NSF had subsidized the cost of registering all domain names. At that time, there were 120,000 registered domain names. In September 1998, when NSF's agreement with NSI expired, the number of registered domain names had passed 2 million.

The year 1998 marked the end of NSF's direct role in the Internet. That year, the network access points and routing arbiter functions were transferred to the commercial sector. And after much debate, the Department of Commerce's National Telecommunications and Information Administration formalized an agreement with the nonprofit Internet Corporation for Assigned Numbers and Names (ICANN) for oversight of domain name registration. Today, anyone can register a domain name through ICANN-accredited registrars.

NSF after NSFNET. The decommissioning of NSFNET and privatization of the Internet did not mark the end of NSF's involvement in networking. NSF continues to support many research projects to develop new networking tools, educational uses of the Internet and network-based applications.

Through its programs, NSF helps research and education institutions—including those serving underrepresented minorities, rural areas and Native American reservations—make and enhance their connections to the Internet. Over the years, NSF has been instrumental in providing international connections services that have bridged the U.S. network infrastructure with countries and regions, including Europe, Mongolia, Africa, Latin America, Russia and the Pacific Rim.

In addition, NSF has continued to extend the reach of the highest-performance U.S. research and education networks by supporting connectivity and collaborations with their counterparts in Canada, Europe and Asia.

Networking for Tomorrow

Credit: Starlight

NSF's history of leadership in cutting-edge network technologies goes as far back as the NSFNET backbone, which catalyzed the growth of today's Internet (*see "A Brief History of NSF and the Internet"*). NSF continues to promote activities to explore and engineer the Internet of the future.

Gigabit Network Testbeds: 1989-1995. In 1989, while the NSFNET backbone was carrying traffic at 1.5 megabits per second, NSF and the Defense Advanced Research Projects Agency (DARPA) provided $20 million for the Gigabit Network Testbed Initiative, which established five testbeds to explore long-distance networking issues and applications at bandwidths a thousand times greater—up to 2.4 gigabits (2,400 megabits) per second.

Commercial network providers and computer companies contributed an estimated $400 million to deploy the testbeds and support their own participation in the research activities. NSF's supercomputing centers in Urbana-Champaign, Ill., Pittsburgh and San Diego helped push the frontiers of high-speed networking in three of the testbeds.

vBNS: 1995-2003. As NSF worked to privatize the mainstream Internet in the mid-1990s, the foundation also signed an agreement with MCI in April 1995 to establish the very-high-performance Backbone Network Service (vBNS). Without competition from general Internet traffic, the vBNS permitted advanced networking research and the development of novel scientific applications. The vBNS initially connected the NSF supercomputer centers and NSF-specified Network Access Points, where linked the vBNS to other research networks.

The vBNS began operating at speeds of 155 megabits per second (Mbps), or OC-3, at a time when the fastest general Internet links operated at 45 Mbps. In 1997, the vBNS backbone was upgraded to 622 Mbps (OC-12). The vBNS also supported new technologies, such as IPv6, to meet the special needs of advanced applications. By 2000, the vBNS backbone was upgraded to 2.4 gigabits per second (OC-48). In 2000, NSF awarded MCI a three-year, no-cost extension to continue operating the vBNS for university customers; commercial connections to vBNS were also offered for the first time.

With the start of the vBNS, NSF began a series of programs designed to help universities and research institutions join advanced high-performance research networks. In its eight-year history, the High-Performance Network Connections program helped 250 institutions enhance their high-end network connectivity.

Advanced Networking Research: 1995-Present. With the start of the vBNS, NSF also initiated a research program to provide technical and engineering support and overall coordination of the vBNS connections. The National Laboratory for Advanced Network Research (NLANR) was created in 1995 as a collaboration among the NSF supercomputer centers. As the vBNS evolved into a stable leading-edge platform and other high-speed networks were formed, NLANR expanded its focus.

Today NLANR offers applications and user support, engineering services and measurement and network analysis to institutions that are qualified to use high-performance network service providers, such as Internet2's Abilene network and STAR TAP (described below).

Next-Generation Internet: 1996-2000. The second half of the 1990s was a period of widespread activity in the area of high-performance networking. In October 1996, the

The United States, Russia and China completed the first round-the-world computer network ring, which will be used for joint scientific and educational projects. In closing the ring, the three nations increased the bandwidth between the United States and China and made the first-ever fiber network connection across the Russia-China border. The image shows the route of the Little GLORIAD network as it passes through Chicago, Amsterdam, Moscow, Novosibirsk, Zabajkal'sk, Manzhouli, Beijing and Hong Kong.

Credit: Trent Schindler, National Science Foundation

President announced the Next-Generation Internet (NGI) initiative and pledged $300 million to connect universities and national laboratories with high-performance networks and to promote next-generation networking technologies.

NGI was not itself an advanced network, although the term is often used generically to refer to future Internet possibilities. However, through the NGI initiative, NSF, the Departments of Defense and Energy, NASA, the National Institutes of Health and the National Institute for Standards and Technology coordinated their advanced networking activities.

NSF's vBNS played a key part in these efforts, and the NGI initiative provided NSF nearly $75 million over three years. This helped 150 institutions get connected to vBNS and later, Internet2's Abilene network, through NSF's High-Performance Network Connections program, exceeding the NGI goal of 100 institutions.

Also as part of NGI, NSF supported the development of hundreds of advanced networking applications, made awards for research on high-performance networking capabilities and established the Science, Technology and Research Transit Access Point (STAR TAP) in Chicago, which connected six U.S. research networks and 12 international research networks by 1999.

The NGI initiative was overseen by the Committee on Computing, Information and Communications (CCIC) through its Large-Scale Network working group, co-chaired by NSF's George Strawn and NASA's Dave Nelson. Although the NGI initiative ended successfully in 2000, the participating agencies continue to interact through the Large-Scale Networking program area of the National Science and Technology Council's Interagency Working Group on Information Technology Research and Development.

Internet2: 1996-Present. While NSF and other federal agencies were deploying their own advanced research and education networks, a group of universities formed the nonporfit Internet2 consortium in 1996 to develop new Internet technologies and capabilities. Today, Internet2 has more than 220 university members; more than 60 corporate sponsors, partners and members; and more than 40 affiliate members including NSF.

The Internet2 consortium also established its own high-performance network named Abilene. The Abilene network began operation in February 1999; a 2.5-gigabits-per-second (Gbps) backbone was completed later that year. Today, Abilene has a 10-Gbps (OC-192c) backbone that connects regional network aggregation points—known as gigaPoPs—to provide advanced network capabilities to Internet2 member institutions in all 50 states, the District of Columbia and Puerto Rico. In 1999, the vBNS and Abilene networks were connected through the NSF-supported STAR TAP.

Many Internet2 member institutions have received NSF High-Performance Network Connections awards. In 1999, NSF made a $6 million award to EDUCAUSE to establish the Advanced Networking Project for Minority Serving Institutions (AN-MSI).

International Bridges: 1990-Present. As far back as 1990, NSF supported programs to link the U.S. and international research and education communities and to assist other countries in connecting to the global Internet. From 1997 to 2004, NSF's High-Performance International Internet Services (HPIIS) program supported STAR TAP in Chicago for interconnecting NSF's vBNS with other advanced networks. In 2002, an NSF award created the optical StarLight exchange point.

The HPIIS project also made awards for sharing some of the costs of the high-performance connections between the United States and the Asia-Pacific region (TransPAC), Russia (MirNET, now NaukaNet) and Europe and Israel (Euro-Link) to the high-performance vBNS and Abilene networks. In addition, NSF is also supporting, through workshops and other awards, a developing project to link Central and South America (AMPATH).

In late 2003, NSF and NaukaNet collaborated with organizations in Russia and China to create the first global-ring network to circle the Northern Hemisphere. Called Little GLORIAD, this network is the first step towards a higher-speed network known simply as GLORIAD, shorthand for Global Ring Network for Advanced Application Development. The ultimate goal of the GLORIAD project is to create a 10 gigabit-per-second optical network around the entire Northern Hemisphere.

That effort got underway in earnest in early 2005, when GLORIAD was one of five projects funded under NSF's new International Research Network Connections program, which continues the HPIIS mission of enabling international research and education collaborations. The other IRNC projects:

- The Consortium of International Research and Education Network (CIREN), which will provide high-performance connectivity between the US and Asia, starting with high-speed connections linking the U.S. West Coast with Tokyo, Hong Kong and Singapore.
- The Western Hemisphere Research and Education Networks (WHREN), which will provide high-performance computing and networking services between North and South America, including connectivity to RedCLARA, Latin America's emerging backbone for research and engineering.

- Translight/Starlight, which will connect exchange points in Chicago, New York and Amsterdam. This award includes funding of a 10 gigabit-per-second connection between the Internet2/Abilene network in the United States, and GEANT, its counterpart network in Europe.
- TransLight/PacificWave, which will connect Seattle and Los Angeles with Australia and other key countries around the Pacific Rim.

A GRAND CONVERGENCE

ESSENTIAL, NOT OPTIONAL

DISCOVERY, LEARNING AND LEADERSHIP

CLASSROOM RESOURCES

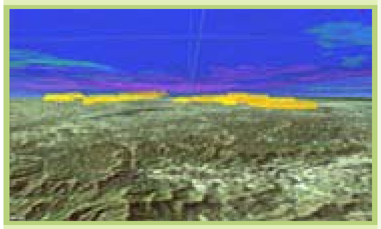

A Georgia Tech weather visualization system displays mesocyclone severe weather cells passing over the north Georgia mountains. Mesocyclone cells are features derived from Doppler radar data, which show the position, motion and extent of severe storms.

Credit: Zachary Wartell

Grand Challenges for Information Technology

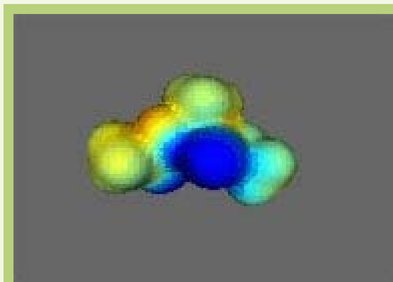

Credit: Curt Breneman, Rensselaer Polytechnic Institute

The needs for and demands placed upon information technology continue to present challenges to the research community. The Computing Research Association's report on the "Grand Research Challenges" in information systems, the result of a three-day workshop supported by NSF, encapsulates the discussions of 65 leading computer science and engineering researchers.

The report identifies five "deliberately monumental" research challenges, each requiring "at least a decade of concentrated research in order to make substantive progress":

- A post-disaster safety net.
- "Cognitive partners" for humans.
- Personalized lifelong learning environments.
- Unfailingly reliable systems.
- Making information technology less complex.

These five grand challenges are, however, only the latest targets in applying information technologies to the most complex scientific and societal issues.

The First Grand Challenges. In 1992 and 1993, NSF, as part of the U.S. High-Performance Computing and Communications (HPCC) program, funded projects by groups pursuing so-called "grand challenges." These grand challenge projects tackled "fundamental problems in science and engineering, with broad economic and scientific impact, whose solution could be advanced by applying high-performance computing techniques and resources." The term "grand challenge," first used by the HPCC program, has been widely adopted in many fields to signify an overarching goal that requires a large-scale, concerted effort.

The HPCC program evolved into today's Information Technology Research and Development (ITRD) interagency effort, which is coordinated by the National Coordinating Office (NCO). NCO supports the President's Information Technology Advisory Committee and manages the High-End Computing Revitalization Task Force. NSF serves as the host agency for the National Coordinating Office.

President's Information Technology Advisory Committee. Established in 1997 and renewed through 2005, the President's Information Technology Advisory Committee (PITAC) was charged with advising the president on high-performance computing and communications, information technology and the Next-Generation Internet.

PITAC's influential 1999 report noted that information technology has created unprecedented possibilities for advancing knowledge across the spectrum of human endeavors, including scientific research, education, engineering design and manufacturing, environmental systems, health care, business, entertainment and government operations. However, the report also asserted that the United States was "gravely underinvesting" in information technology research.

Information Technology Research: 2000-2004. In response to the PITAC report, NSF and other federal agencies worked with the National Science and Technology Council to develop research plans, which were formalized at NSF in the five-year Information Technology Research (ITR) priority area. ITR also built upon goals and ideas from NSF's 1998-1999 Knowledge and Distributed Intelligence (KDI) initiative.

The ITR priority area was designed to encourage and stimulate innovative, high-risk and highreturn research that extended the frontiers of information technology, improved understanding of its impacts on society, helped prepare Americans for the Information Age and reduced the vulnerabilities of society to catastrophic events, whether natural or man-made. The ITR priority area has supported more than a thousand projects that are leading to major advances, new and unanticipated technologies, revolutionary applications and new ways to perform important activities.

Notably, ITR awards have launched a number of community-driven applications of cyberinfrastructure, including the Grid Physics Network (GriPhyN), the National Virtual Observatory (NVO) and the Geosciences Network (GEON). Other ITR-supported research covers topics as diverse as mobile sensor webs for polar research, robotic assistants for the elderly, an International Children's Digital Library, quantum and molecular computing and virtual scientific instruments.

In 2000, the initiative's first year, the NSF ITR priority area focused on projects stressing fundamental research and education. In 2001, applications in science were added. In 2002, the program made awards to support research to create and use cutting-edge cyberinfrastructure, focusing on emerging opportunities at the interfaces between information technologies and other disciplines. By 2003, the ITR budget had grown to nearly $300 million. Reflecting the current world situation, more than 800 proposals received in 2003 were related to homeland security applications.

2004 is the final year of ITR as an NSF priority area and will be a transition year in which changes are made to focus the research and move toward the future. The success of the ITR initiative has ensured that ITR in some form will remain an important part of activities for the Computer and Information Sciences and Engineering (CISE) directorate.

Digital Libraries—Access to Human Knowledge

According to the University of California, Berkeley, report "How Much Information, 2003?", humanity produces about 250 megabytes of data—roughly the text in 250 books—each year for every person on the planet. Only 0.003 percent of this annual output is in printed form, most is in the form of images, sound and numeric data, with more than 90 percent stored digitally.

A participant at the Pacific Educational Conference in Pohnpei, the largest island in the Federated States of Micronesia, undertakes an activity on weaving and symmetry available through the NSF-supported Ethnomathematics Digital Library.

Credit: Nancy Lane, PREL

In the face of this data onslaught, the notion of a "digital library" is a metaphor for thinking about data collections in a networked world. In an effort to preserve data over time and help transform raw data into information and knowledge, digital library research proceeds on three fronts. Meaningful collections from all facets of society must be compiled, structured and preserved. Increasing computational power and network bandwidth must be applied to make these collections accessible, usable and interoperable, and interfaces to these collections must be designed to be clear, flexible and scalable.

DLI and DLI-2: 1994-2003. From 1994 to 1998, NSF, DARPA and NASA funded six digital library projects in the $30 million Phase 1 of the Digital Libraries Initiative. In 1999, NSF, DARPA, the National Library of Medicine, the Library of Congress, NASA and the National Endowment for the Humanities, with participation from the National Archives and the Smithsonian Institution, provided $55 million for Phase 2 (DLI-2). DLI-2 funded 36 projects to extend and develop innovative digital library technologies and applications.

The eSkeletons digital library lets students examine and compare the skeletons of a human, gorilla, baboon and other species through high-resolution annotated photos and 3-D digital models. The collection will soon provide data in a format that will allow users to build replicas of the bones. This image shows a baboon skull with the frontal bone highlighted and a 3-D model of the skull.

Credit: eSkeletons Project, University of Texas, Austin

Today, NSF continues to support digital libraries research in programs through several directorates. DLI-2 and an International Digital Libraries Collaborative Research program are administered within the directorate for Computer and Information Science and Engineering (CISE). NSF's Information Technology Research program also supports several digital library research projects. NSF's directorate for Education and Human Resources (EHR) administers the National Science Digital Library (NSDL), which builds on earlier DLI-2 projects and aims to establish a network of learning environments and resources for science, technology, engineering and mathematics education.

NSF's Geosciences and EHR directorates administer the Digital Library for Earth System Education (DLESE). DLESE is a grassroots effort, affiliated with NSDL, that involves teachers, students and scientists. DLESE encompasses electronic materials such as lesson plans, maps, images, data sets, visualizations, assessment activities, curriculum and online courses.

From Cuneiform Tablets to Google

Many smaller NSF-supported digital library projects—ranging from archaeology to oral history—add value to their collections through services that allow scientists, teachers and students to access, explore, search and interact with the data, as well as connect to information in other collections. They add context to data that might otherwise languish as disconnected content.

The Origins of Google. In 1994, one of the first six NSF digital library projects was awarded to Stanford University and led by Hector Garcia-Molina and Terry Winograd. Two of the graduate students supported by this project—Larry Page and Sergey Brin (funded through an NSF graduate fellowship)—began to explore using the linkages between Web pages as a ranking method for what would become the Google search engine.

International Children's Digital Library. Funded in part by a five-year, $3 million NSF award, the University of Maryland, The Internet Archive, and their partners plan to build a library of 10,000 children's books from 100 cultures as part of a research project to develop new technology to serve young readers. The International Children's Digital Library will serve children aged 3 to 13 years worldwide.

***Informedia Video Digital Library*.** Initiated in 1994, this project at Carnegie Mellon University has pioneered new approaches for full-content search and retrieval of TV and radio news and documentary broadcasts. The current library consists of a 1,500-hour, one-terabyte library of daily news captured over the past two years and documentaries produced for public television and government agencies.

Archeological Sites and Relics. Researchers at UCLA are putting online a real-time computer model of the Roman Forum as it appeared in late antiquity. Another group at UCLA, with collaborators at the Max Planck Institute for the History of Science, in Berlin, is providing scholars with access to a database of tens of thousands of cuneiform tablets and inscriptions.

***National Gallery of the Spoken Word*.** Michigan State University and partner organizations are creating a spoken word collection spanning the 20th century. From Thomas Edison's first cylinder recordings to Neil Armstrong's "one small step for man," this project is addressing challenges such as digital watermarking, compression strategies, audio searches, and user-selectable speech enhancement.

E-Specimens. Two projects at the University of Texas at Austin provide access to libraries of biological specimens. The DigiMorph library is an archive of X-ray computed tomography of biological specimens. The site's animations and details are in use in classrooms and research labs around the world. The e-Skeletons Project lets users examine the bones of a human, gorilla and baboon and information about them in an osteology database.

New Forms of Access. The University of California, Berkeley, Digital Library and the Alexandria Digital Library at the University of California, Santa Barbara, both established in 1994, are researching how digital formats open up new ways to access data. At Berkeley, researchers are exploring how the digital form expands the definition of "documents," especially scholarly publications. The Alexandria Digital Library is building a map collection along with novel tools to access geographically referenced data.

Dolphin Photo-Identification. DARWIN, a system developed by undergraduates at Eckerd College in Florida, assists marine scientists in the study of bottlenose dolphins. The software provides access to a collection of dorsal fin images along with information and sighting data on individuals.

Among other features, users may query the system with an image of an unidentified dolphin's dorsal fin.

A GRAND CONVERGENCE

ESSENTIAL, NOT OPTIONAL

DISCOVERY, LEARNING AND LEADERSHIP

CLASSROOM RESOURCES

Researchers in the United States and Taiwan are studying lake metabolism in a first-of-its-kind project using wireless field sensors on two sides of the Pacific Ocean. Ecologists and computer scientists have deployed sensor-laden buoys in several lakes at the Northern Temperate Lakes field station in northern Wisconsin and in Yuan-Yang Lake in Taiwan (shown here) to measure gross primary production, respiration and net ecosystem production. Wireless networking allows communication with the buoys from anywhere on the Internet, and grid middleware processes the raw sensor data to support cross-site studies and collaborations.

Credit: Lake Metabolism Project

NMI In the Middle

"Middleware" is the glue that holds cyberinfrastructure together by allowing two or more otherwise separate applications to access diverse resources across the network. These advanced network services allow scientists and researchers to collaborate with their colleagues and effectively share instruments, computing resources, laboratories and data.

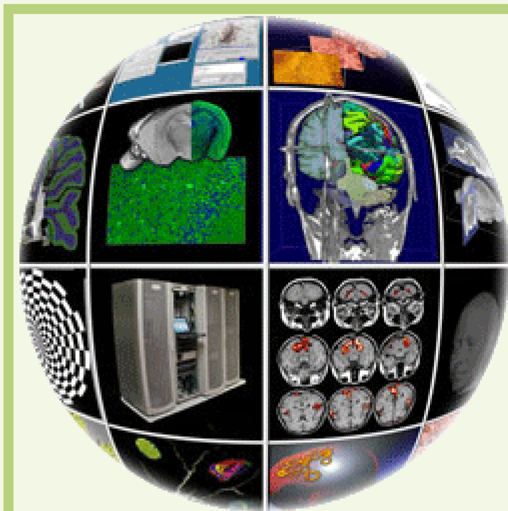

Credit: Biomedical Informatics Research Network

NSF established the NSF Middleware Initiative (NMI) in 2001 to define, develop and support an integrated national middleware infrastructure. NMI activities are making it possible to share scientific resources ranging from telescopes, supercomputing systems and linear accelerators to databases, directories and calendars.

Integration and Standards. In 2001, NSF launched NMI with awards totaling almost $12 million. These awards formed the Grid Research Integration Deployment and Support (GRIDS) Center and the Enterprise and Desktop Integration Technologies (EDIT) Consortium. The GRIDS Center and EDIT Consortium have worked closely with the NSF Partnerships for Advanced Computational Infrastructure (PACI) and private industry to define, develop and deploy an integrated national middleware infrastructure.

Thus far, GRIDS Center and EDIT Consortium have issued four releases of production-quality, open-source and open-standards middleware tools. Available free to the public from the NMI Web site, the components fill functions needed by the research and education community in such areas as user authentication and authorization, resource identification and allocation, job management and scheduling.

Portals and Instruments. NMI awards in 2003 extended the initiative's efforts in production-quality middleware to grid portals and integration of scientific instrumentation. The Open Grid Computing Environment (OGCE) team is focused on simplifying the development of "grid portals," Web-based user interfaces to applications that may access a broad array of resources and services on the grid. The Common Instrument Middleware Architecture (CIMA) team is developing standard grid middleware architecture for access and integration of scientific instruments.

Testbed and Experimental Capabilities. One important emphasis for NMI is to explore ways in which grid computing can be integrated with enterprise computing on university campuses. To that end, NSF has funded the NMI Integration Testbed, consisting of eight universities that coordinate closely to deploy and evaluate NMI middleware in production environments.

The testbed sites use and evaluate software, services and architectures that make it possible for faculty and campus projects to access distributed electronic resources. Testbed efforts gauge the middleware's practicality, emphasizing factors such as performance, ease of use and technical support.

Aspart of NMI, NSF has also made awards that focus on experimental applications of new middleware capabilities as well as near-term additions to the portfolio of middleware services of the NMI releases. These NMI awards include efforts to develop collaboration tools, scalable video services, essential software libraries for grid-based parallel computing, generalized grid-based application environments and tools for grid-based databases.

The goal for these awards is to develop middleware standards and services that will eventually be integrated into the production NMI releases. One of these awards, for example, led to the new H.350 standard for videoconferencing middleware recently adopted by the International Telecommunications Union.