APPENDIX

AGENCY ROLES IN THE INFORMATION TECHNOLOGY FOR THE TWENTY-FIRST CENTURY (IT^2) INITITATIVE

Table of Contents

Department of Defense (DoD)

The Department of Defense's (DoD) participation in this initiative will be supervised by the Under Secretary of Defense for Acquisition and Technology in his role as principal staff assistant to the Secretary of Defense for research and development programs within the DoD. DoD's activities will involve support for basic long-range research and related equipment necessary to facilitate advances in information technology managed through the University Research Initiative program, focused research activities to be managed by the Defense Advanced Research Projects Agency (DARPA), and a new Advanced Research and Development Activity (ARDA) in Information Technologies in support of the Intelligence and Information Security Communities.

The Initiative is Critical to DoD

This initiative will attack problems that will enable revolutionary capabilities of importance to the DoD and the nation as a whole. DoD's priorities in this regard are aligned with Joint Vision 2010, the Chairman of the Joint Chiefs' conceptual template for achieving new levels of warfighting effectiveness. This vision forecasts dynamic change in the nature of potential adversaries, emphasizes the increasingly critical nature of technological advances and their implications, and outlines the emerging importance of information superiority. In short, the DoD is participating because the military that stays ahead in information and its timely dissemination is likely to dominate in the future.

University Initiatives

The University Research Initiative is an established, well-regarded DoD program designed to enhance universities' capabilities to perform basic science and engineering research and related education in science and engineering areas critical to national defense. The Information Technology Initiative will include support through the University Research Initiative focused on stimulating fundamental advances in information science and engineering areas of interest to DoD. DoD plans to support multi-disciplinary research and the purchase of equipment necessary to perform fundamental research work that relates to information science and technology.

DARPA Focused Research Programs

I. Software Research

A particular area of concern to the DoD is the degree to which today's interactive software, much of which was pioneered by DARPA-sponsored researchers, requires the constant attention of human users, who both supply the inputs and consume the results. Although this interactive approach to computation seemed appropriate during the transition from mainframes to personal computers, we must now lay the software foundations for a new era in which our computers will vastly outnumber our population. To accommodate this next step forward, we must develop new approaches to software that will allow individual users to leverage hundreds, and eventually thousands, of networked processors. Such new software approaches include active software that is capable of deploying itself to nodes within the network and autonomous software that allows robots (and knowbots) to react to changes in their environment, thereby freeing humans from the need to painstakingly monitor the actions of every device and agent under their control. In addition, tolerant software technologies will allow systems built from parts that are imperfect to deliver answers that are good enough and given soon enough, as opposed to perfect answers provided too late.

II. Human-Computer Interaction and Information Management

DoD's efforts to free our nation's soldier and citizenry from the tyranny of the keyboard will involve the development of critical technologies to enable more effective interaction between human users and the vast data of future information systems. Technologies to be developed include multi-modal interfaces and the ability to access information regardless of the language the data was originally entered in.

III. Scalable Networks

Current Internet technology targets only two percent of all computers, i.e., the personal computers, servers and supercomputers that are directly responsive to human operators. The remaining 98 percent of computers are stranded within devices whose sensors and actuators are in direct contact with the physical world. Imagine the productivity gains that could be achieved if we were to extend the "depth" of the network to reach these embedded devices and close the chasm between the virtual world, of networked bits and bytes, and the physical world, which we strive to monitor and manage. DARPA will develop methods for accessing, tasking, and managing dynamic, multi-node networks of sensors and fusing and analyzing the resulting data. DARPA will also focus on the network technology required to enable these dynamic networks of sensors, actuators, processors and data storage devices.

IV. High End Computing

DoD's focus at the high end will be on the development of component technologies that, while suitable for use in next generation supercomputers, also have the potential to revolutionize the design of systems that acquire and process information that is directly acquired from the physical world. These include new approaches to computer component technologies and system architectures.

DARPA's Role

DARPA has a tradition of deliberate entrepreneurship and risk-taking to achieve high payoff goals — especially in areas where the level of technical risk would inhibit others. The focused research areas of DoD's Information Technology Initiative participation are aligned with DARPA's investment strategy, which: (1) concentrates on high-payoff technologies and concepts; (2) sometimes invest in those critical technology developments, such as the ARPANET and internetworking technology, that require an extended focus and long-term development; (3) is idea-driven and outcome oriented, i.e., creates a vision of what could be done in 10 to 20 years to solve a specific problem and defines the technology development that must be pursued to make the vision a reality. DARPA will channel DoD resources into selected technologies and pursue them with vigor and determination. Just as critical, however, are the basic research programs of DoD and elsewhere in the Federal Government that sustain the research base on which DARPA's efforts must be built.

ARDA

The Defense Department and the Intelligence Community are also launching a program known as ARDA, the Advanced Research and Development Activity in Information Technology for the Intelligence and Information Security Communities. ARDA was created because both of these communities believe that they are currently underfunding high-risk/high-payoff research in information technology. In addition to supporting fundamental research and development in key areas of information technology, ARDA will also work to recruit “research fellows” to supplement the Intelligence Community’s R&D workforce, transfer technology that is developed or discovered as a result of the initiative, and form strategic partnerships with world-class researchers in industry and academia.

ARDA will focus its long-term research on a small number of problems that are relevant to the Information Security Community, the Intelligence Community, or support enabling technologies that are relevant to both communities. Examples of potential research areas and objectives include, but are not limited to:

A security management infrastructure that protects data from both active attack or from unintended, passive access;

Securing Intelligence Community and Information Security Community systems and communications as they become increasingly integrated into the Global Information Infrastructure;

Screening, filtering, identifying, categorizing, mining, clustering, selecting, and retrieving multimedia data and documents based on the information needs of intelligence analysts and customers of intelligence;

Automating (partially or fully) selected human analytic processes and cognitive activities;

Enabling information technologies such as advanced computing, mass storage, networking, microelectronics, reverse engineering, mathematics, cryptology, and computational linguistics; and

Methodologies for testing and evaluating information technology research results for efficiency, accuracy, flexibility, portability, usability and scalability across a wide problem set.

Department of Energy's (DOE)

Background. DOE is a science-based agency that contributes to the future of the nation by ensuring our energy security, maintaining the safety and reliability of our nuclear stockpile, cleaning up the environment from the legacy of the Cold War, and enhancing our understanding of the physical world through research in fundamental science.

To accomplish its mission objectives and continue to effectively solve challenging problems of national importance in the 21st century, the DOE has proposed an investment to take it into a new era of computational scientific simulation—to rapidly develop and deploy a new generation of computational simulation capabilities along with the supporting national terascale computing infrastructure and to apply this new capability to large-scale civilian science and engineering problems. This capability will not only revolutionize the Department's approach to solving the most demanding, mission-critical problems, but will stimulate our national system of innovation.

Computer-based scientific simulation is one of the most significant developments in the practice of scientific enquiry in the 20th Century. Computational simulation has dramatically advanced our understanding of the fundamental properties of matter (e.g., the structure and interactions of molecules) and is beginning to have a major impact on critical engineering problems (e.g., automobile design). In many cases, simulation is the only economically and technically viable means to accomplish the DOE missions due to the complexity of the problems and the expense or duration of the needed experiments and tests. In some instances, simulation is the only solution because it is not possible to carry out the needed experiments or tests.

As an example, in 1996, the Clinton Administration launched DOE's Accelerated Strategic Computing Initiative. This program fundamentally changes the way our nation ensures the safety, reliability, and performance of our nuclear weapons stockpile by shifting from an emphasis on nuclear tests to one based on computational simulation validated by experiment.

DOE Role. DOE's participation in the President's IT^2 Initiative is built on a long and successful history as a lead agency in computational science, computer science, applied mathematics, and high-performance computing; in the design, development and management of large-scale user-facilities, including a number of national computing facilities; and in the development of cross-cutting computing and engineering technologies such as data analysis and visualization, collaborative environments, and other related technologies.

DOE's unique set of research capabilities, mission goals, and high-performance computing and advanced applications expertise are critical to three primary goals of the President's IT^2 Initiative:

GOAL #1 - SCIENTIFIC APPLICATIONS. To design, develop, and deploy computational simulation capabilities to solve scientific and engineering problems of extraordinary complexity.

GOAL #2 - COMPUTER SCIENCE AND ENABLING TECHNOLOGY. To discover, develop, and deploy crosscutting computer science, applied mathematics, and other enabling technologies.

GOAL #3 - SCIENTIFIC SIMULATION INFRASTRUCTURE. To establish a national terascale distributed scientific simulation infrastructure.

Goal #1. SCIENTIFIC APPLICATIONS. DOE will initially focus on two major simulation projects that are critical to the agency's mission, have urgent deadlines, are of high scientific impact, and are well positioned to take advantage of terascale computing—global systems and combustion systems. In addition, a number of applications in basic science (e.g., genomics) that fulfill these criteria will be initiated.

Applications	Project Goal	Benefits
Atomic, molecular, gas/fluid, and device-level simulation of chemical reactions, fluid dynamics, and energy transfer in real-world **combustion devices**	To enable **predictive** simulation of complex, practical combustion systems that underlie 80% of global energy usage.	Design of improved combustion devices (e.g., automobile engines) leading to: • Reduced pollution–close to zero for nitrogen oxides and soot. • Improved efficiency–reduces fuel consumption, minimizes carbon dioxide emissions
Advanced modeling of the impact of human activities on **global systems** such as weather patterns and long-term climate change	To predict, by the year 2005, regional climate under scenarios of global energy usage, in a multi-agency, international effort to support the IPCC 4th Assessment.	• Improved management of water resources • Estimates of regional energy demand and use • Anticipating damage caused by natural disasters
Basic sciences including, but not limited to, material science, structural genomics, plasma physics, high energy physics, and subsurface transport	Project goals are specific to the research areas (to be selected by peer review), but will in all cases dramatically alter research progress though revolutionary applications of simulation.	• New concepts for and understanding of scientific problems of national importance that cannot be obtained from theory or experiment alone • Accelerated development of new energy technologies from knowledge gained in scientific simulations

Goal #2. COMPUTER SCIENCE AND ENABLING TECHNOLOGY. DOE will focus on the development of the computer science and applied mathematics technologies that enable the advances noted above in the scientific applications. Critical to the success of the applications is the development and deployment of advanced technology in computational algorithms and methods, and software libraries; problem solving and code development environments and tools; distributed computing and collaborative environments; visualization and data management systems; and computer systems architecture and hardware strategies.

Goal #3. SCIENTIFIC SIMULATION INFRASTRUCTURE. The hardware strategy will be driven by the applications requirements and will be based on the acquisition of a balanced system of advanced computers for computational methods and software development as well as the demanding applications listed in GOAL #1. Major computing platforms are planned for acquisition by DOE with an aggregate of 5 teraflops in FY 2000 and 40 teraflops by the year 2003. Local random-access memory, on-line disk, and archival storage needs will increase from 1-2 terabytes, 30-40 terabytes, and approximately 5 petabytes, respectively, in FY 2000 to 10-15 terabytes, 300-500 terabytes, and 15-20 petabytes, respectively, in FY 2003. The greatly expanding computational capacity of these resources coupled with the distributed nature of the scientific teams collaborating on the applications requires advanced communications capabilities to support distributed collaborations, distributed simulations, remote visualizations, remote steering of computations, and the sharing of huge datasets.

While the DOE's major contribution to the President's IT^2 Initiative relates to the advanced computation and engineering component, its efforts in crosscutting and enabling technologies are expected to both drive and profit from basic information technology research.

National Aeronautics and Space Administration (NASA)

Background: NASA is an investment in America's future. As explorers, pioneers, and innovators, we expand frontiers in air and space to inspire and serve America and to benefit the quality of life on Earth. NASA's bold missions in space and aeronautics require significant advances in information technology. Potential benefits of NASA's participation in the Information Technology for the Twenty-first Century Initiative are:

- To develop **autonomous spacecraft and rovers**, that will enable us to explore the universe with less risk and at lower cost**.**
- To increase our ability to transform large and distributed streams of data into **scientific data understanding and knowledge**.
- To promote new concepts in **aviation operations** that will enable aircraft to safely and effectively negotiate clearances, routings and sequencing.
- By enabling **an intelligent synthesis environment** that will revolutionize the nation's science and engineering infrastructure – the tools, technologies, capabilities, and work practices we use to deliver cutting-edge capabilities

NASA's participation in the President's initiative is built on a long and successful history in computational science, computer science, and high-performance computing. NASA is an innovator and early user of information technologies. NASA's unique mission and research capabilities will be applied to support the three goals of the President's IT^2 initiative:

Goal #1 - Long-term information technology. To discover, develop and deploy fundamental computer science, computational science and applied mathematics to serve the national economy, national security, education, the global environment and broad societal goals.

Goal #2 - Advanced computing for science, engineering, and the Nation. Empower computational discovery in all areas and ensure tera-scale computational resources are available to a broad research community to enable scientists and engineers to solve problems of extraordinary complexity.

Goal #3 - Training of the information technology workforce. Understand the impact of social, ethical, economic, political, and legal factors on the development of information technology and vice versa in order to mitigate negative socio-economic impacts.

NASA Role Goal # 1. In FY 2000, the first year of a planned presidential initiative, NASA will support fundamental intelligent systems research focused on several areas, including:

Software research – NASA will conduct research in automated reasoning which requires advances in model-based reasoning, high assurance software, biologically motivated adaptive systems, and planning and scheduling. The goal is to demonstrate coordinated planning and execution of self-tasking and self-repairable robotic networks early in the next decade to support future decisions on establishing permanent, sustainable robotic outposts at key points of scientific interest in the solar system. Developed in close coordination with the research objectives of ongoing and future programs, these revolutionary capabilities will dramatically increase the return of NASA's space and Earth science missions.

Human-centered computing and information management – NASA will develop intelligent systems for data understanding including geographically distributed computing, reconfigurable computer architectures, biologically-motivated hardware and software, and knowledge discovery and data mining. NASA will advance human-centered computing including knowledge management and institutional knowledge capture, immersive/haptic environments, internet-based knowledge representation and cognitive architectures. The goal is to provide the technology to understand and manage large and distributed streams of data to support applications such as Earth system science.

High-end computing – NASA will explore revolutionary computing concepts including quantum mechanical computing, neurally-inspired computing, holographic memory devices, optoelectronic systems, biological computing, biomimetics and ubiquitous intelligence.

NASA Role Goal #2. NASA will support relevant research in applications and use of tera-scale infrastructure in implementing this initiative.

Applications – NASA will advance the Intelligent Synthesis Environment, which includes rapid synthesis and simulation on tera-scale systems and distributed collaborative engineering capability. The goal is to demonstrate design and synthesis of vehicles and missions with greatly reduced redesign and rework and improved operational reliability. The first Intelligent Synthesis Environment testbed will apply these benefits to the design of future reusable launch vehicles that promise to greatly reduce the cost of space transportation. By 2003, substantial improvement will be made towards the following objectives as measured against a standard vehicle and mission baseline: reduce design and development time to 12 to 18 months, reduce testing requirements by 75%, predict life-cycle cost to within 10%, and predict life-cycle risks with a confidence level of 95%.

Access to tera-scale infrastructure - NASA applications which require tera-scale capability include understanding seasonal-to-interannual climate change and demonstration of intelligent synthesis environment.

NASA Role Goal #3. NASA will support training and education efforts to revolutionize cultural change in the scientific and engineering community to exploit the power of an intelligent synthesis environment.

National Institutes of Health
(NIH)

The NIH mission is to uncover new knowledge that will lead to better health for everyone. NIH works toward that mission by conducting research in its own laboratories; supporting the research of non-Federal scientists in universities, medical schools, hospitals, and research institutions throughout the country and abroad; helping in the training of research investigators; and fostering communication of biomedical information. One of the world's foremost leaders in medical research, NIH funds almost 30 percent of the Nation's health-related research and development.

To achieve its mission, the NIH relies on bioinformatics, high-end computer systems, high-speed networks, and other advanced information technologies to provide the performance needed in today's biomedical computing and information rich environment. The information generated by fundamental science, such as genetic sequencing and mapping, and the information and computation required for optimal practice of medicine have exceeded the capacity of earlier methods of data storage and manipulation. Further advances in biology, and our ability to diagnose and treat many of our most devastating illnesses and disabilities, will be dependent upon information technology.

The reliance of the medical sciences on the application of information technology can be exemplified by several recent advances and current projects:

The full blueprint for a multi-celled organism *C. elegans*, or roundworm, was announced in December, 1998. This marked a historic step for the Human Genome Project as it is the first time scientists have deciphered the genetic instruction of a complete organism that shares many of the same basic life functions as humans. Because *C. elegans* shares 40% of its genes with humans, it can serve as a model for human genetic research and provide entrée into the discovery of cures to human diseases.

Many of today's most effective interventions are the direct result of knowledge gained through clinical trials–studies that evaluate the safety and effectiveness of new drugs and other interventions. NIH is creating an internet accessible database that will include all federally and privately funded clinical trials for drugs for serious or life threatening diseases and conditions submitted under Investigational New Drug applications.

NIH's support of IT^2 will further the Nation's biomedical research progress, and improve health, through a number of initiatives:

Software and Algorithm Research and Development. The increased complexity and interdependency of biomedical research and computation has created a tremendous need for robust and improved software. There are great requirements in bioinformatics, imaging and statistical analysis as well as development of methods in ontologies (knowledge bases) and three-dimensional image labeling and searching. NIH proposes to:

- Establish programs that will develop software and algorithms designed to meet the needs of biomedical researchers. These programs, in part driven by the rapid expansion of sequencing the human genome, would also create transferable software that NIH-funded researchers can use at national centers and at their own institutions.

- Develop simulation methods for implementing realistic models of molecular, cellular, organ and epidemiological systems on parallel architectures and create component libraries that will provide a framework so that researchers can assemble these parts into applications that serve their specific needs.
- Further the work of the Brain Molecular Anatomy Project (BMAP), a multi-institute initiative that supports research on the genomics of the nervous system, with initial efforts focusing on the discovery of new genes and the study of gene expression. This initiative will provide the capability to quantify and track the expression of tens of thousands of genes in space and time, and will generate enormous amounts of data.

High-End Computing. Supercomputing technology is used across the government for applications such as imaging, and processing the volumes of data such as that generated by the Human Genome Project. To permit labs to establish high-throughput/low-cost processor clusters, NIH proposes to:

- Develop an initiative to permit small labs to establish their own implementations of the NASA-developed “Beowulf" technology of clustered low-cost processors. The strategy is to build a supercomputer using many commercial, off-the-shelf, processors instead of expensive "supercomputer" processors. This funding will provide the kernel for enabling development, software conversion and the creation of a cookbook approach for other labs implementing these systems.

IT Workforce Training. NIH continues to be a leader in the training of biomedical and behavioral researchers. The need to attract cross-disciplinary researchers from information technology into biology is imperative as biology increases to become an information-based science. The paucity of researchers with abilities in both biology and information technology will soon become an impediment in the advancement of scientific research. NIH proposes to:

- Stimulate the integration of individuals with quantitative expertise, such as physicists, engineers, mathematicians and computer scientists, into the biological and imaging sciences. Career development awards would be expanded to include programs to provide the requisite skills to investigators through cross-disciplinary training. The "Individual Mentored Research Scientist Development Award in Genomic Research and Analysis" is an excellent example. The development and successful introduction of new technologies require both an understanding of biology and the fundamentals of the technology to be applied, so that the solutions derived are appropriate and can be implemented.

National Oceanic and Atmospheric Administration (NOAA)

The **National Oceanic and Atmospheric Administration's mission** is to describe and predict changes in the Earth's environment, and to conserve and manage wisely the Nation's coastal and marine resources to ensure sustainable economic opportunities. To support this mission, NOAA develops and uses complex computer models of the atmosphere and oceans, maintains databases approaching a petabyte in size, collects environmental data in real-time both routinely and in response to disasters, and disseminates information products to emergency managers, business and the public for the protection of life and property.

A fundamental approach that NOAA has used over its entire history is computer simulation of the atmosphere and oceans, treated either separately or as interacting systems. Computer modeling is central to its mission, both for weather forecasting and climate research. Weather forecasting requires rapid and robust input of observational data via data collection and communications from sources around the globe, from satellites, and, in the near future, from high density Doppler radar data. These data are assimilated, using advanced variational techniques, into weather models with increasingly sophisticated physics and higher spatial and temporal resolution. A fundamental requirement of weather forecast models is that they produce forecast results within extremely strict wall-clock constraints. This can only be accomplished with the fastest computers available using the most efficient software techniques and algorithms. Climate modeling by its nature requires long running model experiments with increasingly complex model physics. At these time scales, the oceans have a profound influence on the atmosphere, requiring fully interacting representation of the atmosphere, oceans, and land-surface hydrology. As climate research advances and begins to resolve regional climate effects, climate model sophistication and complexity will increase dramatically, not only to address the demands of higher model resolutions, but also to support the demands of more elaborate representations of model physics that these higher resolutions require. Finally, NOAA must be more effective in providing forecast and research information to its many customers. Advanced communications technologies are providing NOAA with new opportunities to enhance this information dissemination, not only within its organization, but also to the public at large in normal times and to emergency managers in times of crisis. This work leverages existing resources to most effectively advance the use of state-of-the-art information technologies to accomplish NOAA's mission.

As NOAA faces new demands for increasingly accurate weather forecasts and more robust predictions of near- and long-term climate, NOAA scientists must address important computer science challenges in the development and implementation of climate/weather applications for scalable architectures. First, as models become increasingly complex, the need for a modular or component-based design becomes more critical, both to improve usability by a broad range of users and to assure maintainability. While NOAA scientists have been incorporating this approach into the design of some of their major climate model codes, further work is needed to refine and test this strategy for its value for collaborative geophysical research. In addition, as NOAA organizations acquire and install highly scalable parallel systems for use in both weather forecasting and climate research, end users are in the process redesigning and optimizing all of their major weather and climate for highly to massively parallel architectures. The acquisition by the National Weather Service of an IBM SP system for its operational weather forecasting mission is an example of this transition and its importance to NOAA's mission. For these transitions to be effective, NOAA will need to develop efficient scalable versions of current NOAA applications, which will require the development of new algorithms and the redesign of model codes to optimize cache use rather than vector performance. The most effective way to achieve this will be for model developers to work side-by-side with computer scientists, both to create effective model designs and to facilitate the development of software tools. Finally, as the resolution and complexity of their models increase, NOAA scientists will require much more effective ways to visualize and therefore better understand the multi-dimensional geophysical data sets that these models produce.

At the beginning of this initiative, NOAA will push the state of the art in the use of advanced high-speed computing, efficient parallel algorithms, and high bandwidth data communications to provide the Nation with the best-available weather forecasts and climate predictions. As the initiative progresses, NOAA will be among the first to adopt technologies to benefit the Nation in large-scale computing, data management for petabyte data systems, data visualization as an aid in the diagnosis of the output of advanced geophysical numerical models, and nomadic real-time streaming of data, sound, and video over robust adaptive networks.

Benefits

Using the results of the FY 2000 IT^2 initiative, over the next 5 to 10 years NOAA expects to realize many critical benefits including the following:

HURRICANES

Reduce 72-hour forecasts of hurricane track error by 20% from 220 to 150 nautical miles.

Recognize hurricane forecast situations that have inherently low predictability through the use of compute-intensive ensemble techniques.

Improve forecasts of hurricane intensity by 20-30% over current forecasts through use of higher resolution models and more sophisticated physics made possible by advanced computing and atmospheric science.

SEASONAL TO INTERANNUAL CLIMATE

NOAA will improve forecasts of seasonal to interannual climate by including the full global air-sea interaction and hydrology, and enhanced probabilistic guidance using ensembles.

TORNADOES

NOAA forecasters expect to be able to forecast tornadoes as much as 2 hours in advance (the current standard is 15 minutes) through the use of ultra-high resolution, limited-area models made possible by advanced information technology as well as advanced atmospheric understanding.

ROUTINE FORECASTS

NOAA will provide 5-day forecasts with the accuracy of current 4-day forecasts, a 20% improvement.

LONG-TERM CLIMATE

NOAA will produce a 50% reduction in current modeling uncertainty through a more complete treatment of clouds. These science-based improvements will produce a more fact-centered basis for optimal policy decisions in the future.

Milestones:

FY 2000

Release two additional features of the Scalable Modeling System, a suite of advanced modules which take advantage of emerging component software technologies to support rapid development and easy porting of atmospheric, climate and ocean models across scalable architectures.

FY2001

Evaluate the capabilities of a more advanced GFDL Hurricane Prediction System for providing improved track forecasts as well as predicting other storm features, such as wind and precipitation fields and changes in storm intensity**.** The system will also be optimized for performance for the Weather Service's scalable architecture using advanced programming and cache management techniques.

FY2002

Demonstrate progress in improving the capabilities of the next-generation GFDL coupled research model for predicting seasonal-interannual climate and for elucidating some of the processes that control El-Niño-Southern-Oscillation events**.** A key aspect of this model development activity will be the use of component-based design to facilitate the exchange of physics packages and maximize the usability of the modeling system.

FY2003

Isolate some sources of climate "drift" and define a strategy for reducing their effect on long-running, higher resolution coupled climate models. A key aspect of the diagnostic process will be robust, flexible data visualization and advanced diagnostic strategies including data mining techniques to locate the sources of drift.

FY2004

Provide higher resolution projections of climate change, with improved representations of clouds and ocean circulation, to the impacts research community as part of the 2005 IPCC climate change assessment. It is expected that these enhancements will be considerably simpler to implement as a result of the component-based design implemented in the early years of the initiative. These complex and subtle research results will be communicated using collaborative visualizations techniques not readily available today.

National Science Foundation
(NSF)

The National Science Foundation funds a very broad spectrum of fundamental research in science, mathematics, and engineering at the Nation's academic institutions. Equally important is the NSF's responsibility for educational activity in the sciences across the full spectrum from K-12 to graduate school to informal education and workforce training. The Foundation's activities in all these areas contribute to, and are aided by, its strong support of research activity in basic information technology, including computer science, computation, and information transfer.

In FY 2000, NSF will support basic efforts focused on several areas of software research, including:

- *Building "no-surprise" performance-engineered software-systems* in a scalable way to deliver functionality, predictability, and security. There will be a strong emphasis on research to extract design principles from large, successful software projects, and to validate promising approaches such as component-based software design.

- *Developing hardware/software co-design to* help build vertically integrated electronic systems that may be contained on a single chip. In the future, such systems will lie at the heart of many critical systems and technologies.

- *Building high-confidence systems* by addressing a range of significant software challenges, such as predictability, reliability, and security.

Efforts in human-computer interactions and information management will focus on research in areas such as:

- *Multiplying individual physical and mental capabilities* via computer sensors and actuators to control devices which aid people at work and at home, and research on high performance communications links, high-end computational engines, and the worldwide information resources to manage them.

- *Meeting, working, and collaborating in cyberspace*, to allow interactions over networks in a realistic, 3-dimensional, virtual environment.

- *Building a ubiquitous content infrastructur*e that enables seamless retrieval of text, data, visual and other available information in all subject areas by all citizens.

In the area of scalable information infrastructure, NSF-supported research will focus on:

- *Broadband tetherless communications* to help realize exciting new technologies such as telemedicine, crisis-management applications, and expanded distance-learning.

- *Understanding, modeling, and predicting the behavior of networks.* Research to simulate and model million-node to billion-node networks, taking into account such tradeoffs as bandwidth, reliability, latency, and different network modalities.

- *Integrating end-to-end performance component*s (networking, data manipulation, algorithms, simulation and modeling, human-computer interfaces, visualization) to achieve overall system performance.

High-end computing research has direct impact on both computer and information science, and areas such as biology, engineering, geosciences, mathematical and physical sciences, and social, behavioral, and economic sciences. Research activities will include:

- *Algorithms related to computational complexity*, particularly algorithms for solving the partial differential equations arising from models of physical phenomena using high-end systems architectures.

- *Access to terascale systems* for computer science researchers, recognizing that many important problems emerge only at large scale. Examples include research on file systems coordinating terabyte data transfers on thousands of disks, and programming systems for managing multilevel memory hierarchies on thousands of processors.

- *Empowering computational discovery* in all of these areas by coupling advanced computation, communications, and data technologies. Research will include work on computational manipulation, modeling, and representation of physical phenomena, as well as on how to achieve the scalability of thousands of processors, million-way parallelism, and distributed and parallel data archives.

Relevant infrastructure is key to implementing this initiative.

- NSF will provide open, competitive access to cutting-edge computing resources for the science and engineering community. Plans are: putting in place a 5-teraflop capability at one site in FY 2000; potentially developing a second site with 5-10 teraflop capability in FY 2001; and providing upgrades. Careful coordination will be undertaken with other agencies to leverage the software, tools and technology investments for mutual benefit. The terascale computing resources will be connected to the PACI network, the core of NSF's computational infrastructure, enabling an acceleration of capabilities planned for the future.

- To make these computing resources available to the broad research community ubiquitous high-bandwidth network connectivity must be available. This will make the location of data and computing resources irrelevant, allowing placement where facilities are available and operating costs are minimized.

There are important ethical, social and information technology workforce issues to confront. These include:

- *Expanding understanding of the impact of social, ethical, economic, political, and legal factors* on the development of information technology and vice versa in order to mitigate negative socio-economic impacts of IT technologies

- *Developing a more skilled American workforce* for the future global marketplace, utilizing the many powerful tools available because of information technologies.