Advanced Scientific Computing Research

Funding Profile by Subprogram

(dollars in thousands)

	FY 2004 Comparable Appropriation	FY 2005 Original Appropriation	FY 2005 Adjustments	FY 2005 Comparable Appropriation	FY 2006 Request
Advanced Scientific Computing Research					
Mathematical, Information, and Computational Sciences	193,879	234,340	-1,872[a]	232,468	207,055
Laboratory Technology Research	2,916	0	0	0	0
Total, Advanced Scientific Computing Research	196,795[b]	234,340	-1,872	232,468	207,055

Public Law Authorizations:
Public Law 95-91, "Department of Energy Organization Act"
Public Law 103-62, "Government Performance and Results Act of 1993"

Mission

The mission of the Advanced Scientific Computing Research (ASCR) program is to deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy. In the past two decades, leadership in scientific computation has become a cornerstone of the Department's strategy to ensure the security of the nation and succeed in its science, energy, environmental quality, and national security missions. According to a number of authorities, ranging from the President's Science Advisor and the President's Council of Advisors on Science and Technology to the National Research Council and the Council on Competitiveness, this scientific leadership should be a priority for the nation.

Benefits

ASCR supports DOE's mission to provide world-class scientific research capacity through peer-reviewed scientific results in mathematics, high performance computing and advanced networks, and through the application of computers capable of trillions of operations per second (terascale computers) to advanced scientific applications. Computer-based simulation enables us to predict the behavior of complex systems that are beyond the reach of our most powerful experimental probes or our most sophisticated theories. Computational modeling has greatly advanced our understanding of fundamental processes of nature, such as fluid flow and turbulence or molecular structure and reactivity. Through modeling and simulation, we will be able to explore the interior of stars to understand how the chemical elements were created and learn how protein machines work inside living cells, which could enable us to design microbes that address critical waste cleanup problems. We can design novel catalysts and high-

[a] Reflects a rescission in accordance with P.L.108-447, the consolidated Appropriations Act, 2005.
[b] Includes a reduction of $1,198,000, for a rescission in accordance with P.L. 108-137, the consolidated Appropriations Act, 2004, a reduction of $4,908,000 which was transferred to the SBIR program, and a reduction of $589,000 which was transferred to the STTR program.

efficiency engines that could expand our economy, lower pollution, and reduce our dependence on foreign oil. Computational science is increasingly central to progress at the frontiers of almost every scientific discipline and to our most challenging feats of engineering.

Strategic and Program Goals

The Department's Strategic Plan identifies four strategic goals (one each for defense, energy, science, and environmental aspects of the mission) plus seven general goals that tie to the strategic goals. The ASCR program supports the following goal:

Science Strategic Goal
General Goal 5, World-Class Scientific Research Capacity: Provide world-class scientific research capacity needed to: ensure the success of Department missions in national and energy security; advance the frontiers of knowledge in physical sciences and areas of biological, medical, environmental, and computational sciences; or provide world-class research facilities for the Nation's science enterprise.

The ASCR program has one program goal which contributes to General Goal 5 in the "goal cascade":

Program Goal 05.23.00.00: Deliver forefront computational and networking capabilities **—** Deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy.

Contribution to Program Goal 05.23.00.00 (Deliver forefront computational and networking capabilities)

The ASCR program contributes to Program Goal 05.23.00.00 by delivering the fundamental mathematical and computer science research that enables the simulation and prediction of complex physical and biological systems, providing the advanced computing capabilities needed by researchers to take advantage of this understanding, and delivering the fundamental networking research and facilities that link scientists across the nation to the computing and experimental facilities and their colleagues to enable scientific discovery. ASCR supports fundamental research in applied mathematics, computer science, computer networks, and tools for electronic collaboration; integrates the results of these basic research efforts into tools and software that can be used by scientists in other disciplines, especially through efforts such as Scientific Discovery through Advanced Computing (SciDAC); and provides the advanced computing and network resources that enable scientists to use these tools to deliver extraordinary science. Applied Mathematics enables scientists to build models of physical and natural systems with extraordinary fidelity, and provides the algorithms the computer requires to manipulate that representation of the world effectively, exposing the underlying structure. Computer science research provides the link between the mathematics and the actual computer systems. Finally, scientific discovery results from simulations conducted on the advanced computers themselves, including experimental computers with hardware designs optimized to enable particular types of scientific applications, and the largest computing capabilities available to the general scientific community. All of these elements are critical to advance the frontiers of simulation. Shrinking the distance between scientists and the resources they need is also critical to the Office of Science (SC). The challenges that SC faces require teams of scientists distributed across the country, as well as the full national portfolio of experimental and computational tools. High performance networks and network research provide the capability to move the millions of gigabytes that these resources generate to the scientists' desktops.

Therefore, the ASCR program contributes to General Goal 5 by enabling research programs across SC, as well as other elements of the Department, to succeed. The following indicators establish specific long term (10 years) goals in Scientific Advancement that the ASCR program is committed to, and progress can be measured against:

- Develop multiscale mathematics, numerical algorithms, and software that enable effective models of systems such as the earth's climate, the behavior of materials, or the behavior of living cells that involve the interaction of complex processes taking place on vastly different time and/or length scales.
- Develop, through the Genomics: GTL partnership with the Biological and Environmental Research (BER) program, the computational science capability to model a complete microbe and a simple microbial community. This capability will provide the science base to enable the development of novel clean-up technologies, bio-energy sources, and technologies for carbon sequestration.

Annual Performance Results and Targets

FY 2001 Results	FY 2002 Results	FY 2003 Results	FY 2004 Results	FY 2005 Targets	FY 2006 Targets
Program Goal 05.23.00.00 Deliver forefront computational and networking capabilities					
Mathematical, Information and Computational Sciences					
		Completed the definitive analysis of the advantages and issues associated with lightweight kernel operating systems rather than full kernels for the compute nodes of extreme-scale scientific computers, resolving a critical issue for the future of high performance computers in the U.S. [Goal Met]			
		Began installation of next generation National Energy Research Scientific Computing Center (NERSC) computer, NERSC-4, that will at least double the capability available in FY 2002 to solve leading edge scientific problems. [Goal Not Met]	Maintained Procurement Baselines. Percentages within (1) original baseline cost for completed procurements of major computer systems or network services, and (2) original performance baseline versus integrated performance over the life of the contracts. [Goal Met]	Maintain Procurement Baselines. Percentages within (1) original baseline cost for completed procurements of major computer systems or network services, and (2) original performance baseline versus integrated performance over the life of the contracts. FY 2005 — <10%	Maintain Procurement Baselines. Percentages within (1) original baseline cost for completed procurements of major computer systems or network services, and (2) original performance baseline versus integrated performance over the life of the contracts. FY 2006 — <10%
		Initiated at least 5 competitively selected interdisciplinary research teams to provide computational science and applied mathematics advances that will accelerate biological discovery in microbial systems and develop the next generation of computational tools required for nanoscale science based on peer review, in partnership with the Biological and Environmental Research (BER) and Basic Energy Sciences (BES) programs, respectively, of submitted proposals. [Goal Met]	*Improved Computational Science Capabilities. Average annual percentage increased in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within the SciDAC effort. [Goal Met]*	*Improve Computational Science Capabilities. Average annual percentage increase in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within the SciDAC effort. FY 2005 — >50%*	*Improve Computational Science Capabilities. Average annual percentage increase in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within the SciDAC effort. FY 2006 — >50%*

FY 2001 Results	FY 2002 Results	FY 2003 Results	FY 2004 Results	FY 2005 Targets	FY 2006 Targets
		Maintained and operated facilities, including NERSC and ESnet, so the unscheduled downtime on average is less than 10% of the total scheduled operating time. [Goal Met]	Focused usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used was accounted for by computations that required at least 1/8 of the total resource. [Goal Not Met]	Focus usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used that is accounted for by computations that require at least 1/8 of the total resource. FY 2005 — 40%	Focus usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used that is accounted for by computations that require at least 1/8 of the total resource. FY 2006 — 40%
Initiated project to understand the advantages and issues associated with lightweight kernel operating systems rather than full kernels for the compute nodes of extreme-scale scientific computers. [Goal Met]	Completed the development of the Cougar lightweight kernel for clusters of Alpha processor-based computers and began the assessment of scalability and performance for selected applications. [Goal Met]				
Continued to fabricate, assemble, and operate premier supercomputer and networking facilities that served researchers at national laboratories, universities and within industry, enabling understanding of complex problems and effective integration of geographically distributed teams in national collaborations. [Goal Met]	Achieved operation of the IBM-SP computer at 5.0 teraflop “peak” performance. These computational resources were integrated by a common high performance file storage system that facilitates interdisciplinary collaborations. Transferred the users with largest data processing and storage needs to the IBM-SP from the previous generation Cray T3E. [Goal Met]				

Means and Strategies

The ASCR program will use various means and strategies to achieve its goals. However, various external factors may impact the ability to achieve these goals.

ASCR will support fundamental, innovative, peer-reviewed research to create new knowledge in areas of advanced computing research that are important to DOE. In addition, ASCR will plan, fabricate, assemble, and operate premier supercomputer and networking facilities that serve researchers at national laboratories, universities, and industry, thus enabling new understanding through analysis, modeling, and simulation for complex problems, and effective integration of geographically distributed teams through national laboratories. Finally, the program will continue its leadership of the SC-wide SciDAC initiative with BES and BER in the areas of nanotechnology and Genomics: GTL. All research projects undergo regular peer review and merit evaluation based on procedures outlined in 10 CFR 605 for the extramural grant program, and under a similar process for the laboratory programs and scientific user facilities. All new projects are selected through peer review and merit evaluation.

External factors that affect the programs and performance include: (1) mission needs as described by the DOE and SC mission statements and strategic plans; (2) evolving scientific opportunities, which sometimes emerge in a way that revolutionizes disciplines; (3) results of external program reviews and international benchmarking activities of entire fields or subfields, such as those performed by the National Academy of Sciences (NAS); (4) unanticipated failures, e.g., in the evaluation of new computer architectures for science, that cannot be mitigated in a timely manner; (5) strategic and programmatic decisions made by other (non-DOE) Federal agencies and by international entities; and (6) the evolution of the commercial market for high performance computing and networking hardware and software.

The fundamental research program and facilities are closely coordinated with the information technology research activities of other Federal Agencies (Defense Advanced Research Projects Administration [DARPA], Environmental Protection Agency [EPA], National Aeronautics and Space Administration [NASA], National Institute of Health [NIH], National Security Agency [NSA], and National Science Foundation [NSF]) through the Computing Information and Communications Research and Development (R&D) subcommittee of the National Science and Technology Council (NSTC), under the auspices of SC and the Office of Science and Technology Policy. This coordination is periodically reviewed by the President's Information Technology Advisory Committee (PITAC). In addition to this interagency coordination, ASCR has a number of partnerships with other programs in SC and other parts of the Department, focused on advanced application testbeds to apply the results of ASCR research. Finally, ASCR has a significant ongoing coordination effort with the National Nuclear Security Administration's (NNSA) Advanced Science Computing (ASC) Campaign to ensure maximum effectiveness of both computational science research efforts.

Validation and Verification

Progress against established plans is evaluated by periodic internal and external performance reviews. These reviews provide an opportunity to verify and validate performance. Quarterly, semiannual, and annual reviews consistent with specific program management plans are held to ensure technical progress, cost and schedule adherence, and responsiveness to program requirements.

Program Assessment Rating Tool (PART)

The Department implemented a tool to evaluate selected programs. PART was developed by the Office of Management and Budget (OMB) to provide a standardized way to assess the effectiveness of the Federal Government's portfolio of programs. The structured framework of the PART provides a means through which programs can assess their activities differently than through traditional reviews. ASCR has incorporated feedback from OMB into the FY 2006 Budget Request and has taken the necessary steps to continue to improve performance.

In the FY 2005 PART review, OMB gave the ASCR program a relatively high score of 84% overall which corresponds to a rating of "Moderately Effective." OMB found that the program supports world-class scientific user facilities, has demonstrated an improved level of interagency communication and cooperation, is in the process of drafting a long-term strategic vision, and has been very successful with a major effort in interdisciplinary software. The assessment found that ASCR has developed a limited number of adequate performance measures which are continued for FY 2006. These measures have been incorporated into this budget request, ASCR grant solicitations, and the performance plans of senior managers. As appropriate, they will be incorporated into the performance-based contracts of Management and Operating (M&O) contractors. To better explain these complex scientific measures, SC has developed a website (http://www.sc.doe.gov/measures/) that answers questions such as "What does this measure mean?" and "Why is it important?" Roadmaps, developed in consultation with the Advanced Scientific Computing Advisory Committee (ASCAC) and also available on the website, will guide reviews every three years by ASCAC of progress toward achieving the long-term Performance Measures. The Annual Performance Targets are tracked through the Department's Joule system and reported in the Department's Annual Performance Report. In response to PART findings, ASCR established a Committee of Visitors (COV) to provide outside expert validation of the program's merit-based review processes for impact on quality, relevance, and performance. ASCR has received the report from the COV and is working on an action plan to respond to the recommendations. In addition, the ASCR strategic plan was approved by the Director, Office of Science, on August 19, 2004. All of these efforts enable ASCR to implement the OMB criteria for basic research programs.

ASCR's role in providing scientific research facilities is strongly supported by the Administration. Funding is provided in FY 2006 to operate the program's facilities at maximum capacity.

Funding by General and Program Goal

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
General Goal 5, World-Class Scientific Research Capacity			
Program Goal 05.23.00.00, Deliver forefront computational and networking capabilities			
Mathematical, Information and Computational Sciences	193,879	232,468	207,055
Laboratory Technology Research	2,916	0	0
Total, General Goal 05, World-Class Scientific Research Capacity	196,795	232,468	207,055

Overview

Computational modeling and simulation are among the most significant developments in the practice of scientific inquiry in the 20th Century. Scientific computing is particularly important for the solution of research problems that are insoluble by traditional theoretical and experimental approaches, hazardous to study in the laboratory, or time-consuming or expensive to solve by traditional means. All of the SC research programs—Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, High Energy Physics, and Nuclear Physics—have identified major scientific challenges that can only be addressed through advances in scientific computing.

ASCR research underpins the efforts of the other programs in SC. The applied mathematics research activity produces the fundamental mathematical methods to model complex physical and biological systems. The computer science research efforts enable scientists to efficiently perform scientific computations on the highest performance computers available and to store, manage, analyze, and visualize the massive amounts of data that result. The networking research activity provides the techniques to link the data producers (e.g., supercomputers and large experimental facilities) with scientists who need access to the data.

ASCR's other principal responsibility is to provide the high-performance computational and networking resources that are required for world leadership in science. Recent dramatic advances in scientific computation by researchers and computer companies underscore the importance of strengthening our position in computational sciences in strategic areas. In March 2002, Japan's NEC Earth Simulator became operational. With a peak speed of 40 teraflops and a demonstrated sustained capability of over 25 teraflops, it is faster by approximately a factor of 50 than the most advanced supercomputer for civilian science—through not for national security work—in the United States. The Administration has recognized the importance of high-end computing. As stated in the "Analytical Perspectives" of the FY 2004 Budget:

> Due to its impact on a wide range of federal agency missions ranging from national security and defense to basic science, high-end computing—or supercomputing—capability is becoming increasingly critical. Through the course of 2003, agencies involved in developing or using high-end computing will be engaged in planning activities to guide future investments in this area, coordinated through the NSTC. The activities will include the development of interagency R&D roadmaps for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues (along with recommendations where applicable) relating to federal procurement of high-end computing systems. The knowledge gained from this process will be used to guide future investments in this area. Research and software to support high-end computing will provide a foundation for future federal R&D by improving the effectiveness of core technologies on which next generation high-end computing systems will rely.

To address these issues the President's Science Advisor chartered the High End Computing Revitalization Task Force (HECRTF), which developed a plan for a Federal research program to address these issues. This task force was co-chaired by SC and the Department of Defense (DOD). The Next Generation Architecture (NGA) activity, begun by ASCR in FY 2004, is one of the key foundations for this effort.

How We Work

The ASCR program uses a variety of mechanisms for conducting, coordinating, and funding research in applied mathematics, network and computer sciences, and in advanced computing software tools. The program is responsible for planning and prioritizing all aspects of supported research, conducting ongoing assessments to ensure a comprehensive and balanced portfolio, regularly seeking advice from stakeholders, supporting core university and national laboratory programs, and maintaining a strong infrastructure to support research in applied mathematics, network and computer science, and advanced computing software tools. The quality of the research supported by the ASCR program is continuously evaluated through the use of merit-based peer review, scientific advisory committees, and interagency coordinating bodies.

Advisory and Consultative Activities

The Advanced Scientific Computing Advisory Committee (ASCAC) provides valuable, independent advice to DOE on a variety of complex scientific and technical issues related to the ASCR program. The ASCAC is charged with providing advice on promising future directions for advanced scientific computing research; strategies to couple advanced scientific computing research to other disciplines; and the relationship of the DOE program to other Federal investments in information technology (IT) research. ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program. This advisory committee plays a key role in assessing the scientific and programmatic merit of presently funded activities and in evaluating plans for the future. The Committee formally reports to the SC Director and includes representatives from universities, national laboratories, and industries who are involved in advanced computing research. Particular attention is paid to obtaining a diverse membership with a balance among scientific disciplines, institutions, and geographic regions. ASCAC operates in accordance with the Federal Advisory Committee Act (FACA), Public Law 92-463; 92nd Congress, H.R. 4383, October 6, 1972, and all applicable FACA Amendments, Federal Regulations, and Executive Orders.

The activities funded by the ASCR program are coordinated with other Federal efforts through the Interagency Principals Group, chaired by the President's Science Advisor, and the Information Technology Working Group (ITWG). The Federal IT R&D agencies have established over a decade of highly successful collaborative accomplishments in multiagency projects and in partnerships with industry and academic researchers. The multiagency approach leverages the expertise and perspectives of scientists and technology users from many agencies who are working on a broad range of IT research questions across the spectrum of human uses of information technology. DOE has been an active participant in these coordination groups and committees since their inception and the ASCR program will continue to coordinate its activities through these mechanisms including an active role in implementing the Federal IT R&D FY 2002-2006 Strategic Plan under the auspices of the NSTC and the President's Science Advisor.

In addition, ASCR, both through ASCAC and independently, supported a number of workshops to support its planning. These include:

- Blueprint for Future Science Middleware and Grid Research and Infrastructure, August 2002 (http://www.nsf-middleware.org/MAGIC/default.htm);
- DOE Science Network Meeting, June 2003 (http://gate.hep.anl.gov/may/ScienceNetworkingWorkshop/);
- DOE Science Computing Conference, June 2003 (http://www.doe-sci-comp.info);
- Science Case for Large Scale Simulation, June 2003 (http://www.pnl.gov/scales/);
- Workshop on the Road Map for the Revitalization of High End Computing (http://www.cra.org/Activities/workshops/nitrd/);
- Cyber infrastructure Report (http://www.cise.nsf.gov/sci/reports/toc.cfm); and
- ASCR Strategic Planning Workshop (http://www.fp-mcs.anl.gov/ascr-july03spw).

Finally, in FY 2003 and FY 2004, ASCR participated in the HECRTF government-wide planning activity, which was established by the President's Science Advisor under the NSTC. The resulting "Federal Plan for High-End Computing, report of the High End-Computing Revitalization Task Force (HECRTF)" (http://www.itrd.gov/pubs/2004_hecrtf/20040702_hecrtf.pdf), which was released on April 13, 2004, plays a critical role in ASCR planning. This plan will be referred to as the "HECRTF Plan."

Facility Operations Reviews

The ASCR program has undertaken a series of operations reviews of the National Energy Research Scientific Computing Center (NERSC), the Energy Sciences Network (ESnet), and the Advanced Computing Research Testbeds (ACRTs).

NERSC, operated by the Lawrence Berkeley National Laboratory (LBNL), annually serves about 2,000 scientists throughout the United States. These researchers work at DOE laboratories, universities, industrial laboratories, and other Federal agencies. Allocations of computer time and archival storage at NERSC are awarded to research groups based on a review of submitted proposals. As proposals are submitted, they are subjected to peer review to evaluate the quality of science, the relevance of the proposed research to SC goals and objectives, and the readiness of the proposed application to fully utilize the computing resources being requested.

The ESnet, managed and operated by the LBNL, is a high-speed network serving thousands of DOE scientists and collaborators worldwide. A pioneer in providing DOE mission-oriented high-bandwidth, reliable connections, ESnet enables researchers at national laboratories, universities, and other institutions to communicate with each other using the leading edge collaborative capabilities, not available in the commercial world, that are needed to address some of the world's most important scientific challenges. The ESnet Steering Committee (ESSC) was established in 1985 to ensure that ESnet meets the needs of SC programs. All program offices in SC appoint members, who represent their scientific communities, to serve on the ESSC. The ESSC is responsible for reviewing and prioritizing network requirements, for establishing performance objectives, and for proposing innovative techniques for enhancing ESnet capabilities. In addition to the ongoing oversight from the ESSC, ASCR conducts external peer reviews of ESnet performance on a three year interval. The last such review was chaired by a member of ASCAC and took place in September 2001. The next review is planned for Spring of 2005.

ACRTs play a critical role in testing and evaluating new computing hardware and software. These evaluations are called Research & Evaluation (R&E) prototypes in the HECRTF Plan. Current R&E prototypes are located at the Oak Ridge National Laboratory's (ORNL's) Center for Computational Sciences (CCS) (Cray X1 Technology and SGI large shared memory technology). In FY 2002, ASCAC conducted a review of NERSC and the ACRTs. The charge to ASCAC, posed the following questions:

- What is the overall quality of these activities relative to the best-in-class in the U.S. and internationally?
- How do these activities relate and contribute to Departmental mission needs?
- How might the roles of these activities evolve to serve the missions of the SC over the next three to five years?

The essential finding of the Subcommittee was that NERSC and the ACRTs are among the best worldwide. It was the opinion of the Subcommittee that these ASCR activities and the related spin-off research efforts contribute significantly to the mission needs of DOE, and profoundly and positively impact high performance computing activities worldwide. The complete report is available on the web. (http://www.science.doe.gov/ascr/ASCAC-sub.doc).

In FY 2001, ASCR conducted a peer review of the CCS at ORNL. The findings from this review validated the contributions that the CCS made to the ACRT activity within the ASCR program. In FY 2004, ASCR conducted a peer review of the CCS evaluation of the Cray X1 computer. The results from this review validated the exceptionally effective results of the evaluation and its contributions to the Federal high performance computing effort. Also in FY 2004, ASCR conducted a peer reviewed competition to establish a Leadership Class Computing facility for Open Science. This competition was won by a partnership of ORNL's CCS with Argonne National Laboratory (ANL) and Pacific Northwest National Laboratory (PNNL) that located the Leadership Class Facility at the CCS.

Program Reviews

The ASCR program conducts frequent and comprehensive evaluations of every component of the program. Results of these evaluations are used to modify program management as appropriate.

In FY 2003, ASCR conducted a peer review of the Numerical Linear Algebra, Optimization, and Predictability Analysis areas within the Applied Mathematics activity. These areas represent 33% of this activity. In FY 2004, ASCR conducted a peer review of the Differential Equations and Advanced Numerical Methods for High Performance Computing areas within the Applied Mathematics activity, representing an additional 33% of this activity. In FY 2005, ASCR will conduct a peer review of the remaining 34% of the Applied Mathematics activity, which consists of Computational Fluid Dynamics and Meshing Techniques. Also, in FY 2003 ASCR completed reviews of all of the SciDAC Integrated Software Infrastructure Centers (ISICs). There are a total of seven such centers (three with a mathematics focus and four with a computer science focus), and this represents over 50% of the ASCR SciDAC budget.

In FY 2003, ASCR also conducted peer reviews of all the SciDAC Collaboratory Pilot and Middleware Projects. These reviews focused on accessing progress and the possible need for mid-course corrections.

In FY 2002, following a comprehensive peer review, the ASCR program approved a proposal from the LBNL to manage and operate the NERSC for FY 2002–FY 2006.

Planning and Priority Setting

The ASCR program must coordinate and prioritize a large number of goals from agency and interagency strategic plans. One of the most important activities of ASCAC is the development of a framework for the coordinated advancement and application of network and computer science and applied mathematics. This framework must be sufficiently flexible to rapidly respond to developments in a fast paced area of research. The key planning elements for this program are:

- The Department and SC Strategic Plan, as updated through program collaborations and joint advisory committee meetings. (http://www.sc.doe.gov/Sub/Mission/mission_strategic.htm);
- SciDAC plan delivered to Congress in March 2000. (http://www.science.doe.gov/scidac/);
- ASCAC report on the Japanese Earth Simulator. (http://www.sc.doe.gov/ascr/ascac_reports.htm);
- The HECRTF Plan (http://www.itrd.gov/pubs/2004_hecrtf/20040702_hecrtf.pdf)

How We Spend Our Budget

The ASCR program budget has two subprograms: Mathematical, Information and Computational Sciences (MICS) and Laboratory Technology Research (LTR). The MICS subprogram has two major components: research and facility operations. The FY 2006 budget request continues the core and SciDAC research efforts and strengthens the research partnerships with other SC offices. Network operations expenditures account for 22% of the national laboratory research. The LTR subprogram was brought to a successful completion in FY 2004.

Advanced Scientific Computing Research Budget Allocation FY 2006

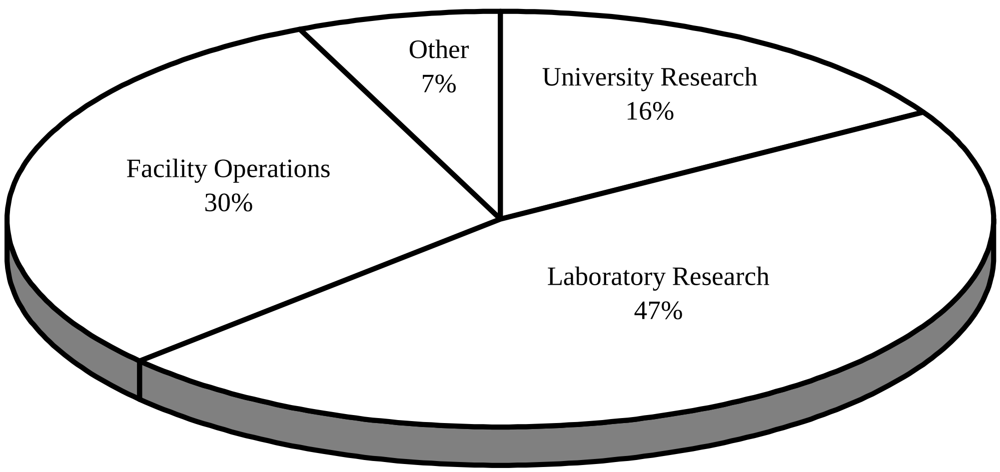

Research

63 percent of the ASCR program's FY 2006 funding will be provided to scientists at universities and laboratories to conceive and carry out their research. National laboratory research scientists work together with the other programs of SC to develop the tools and techniques that allow those programs to take advantage of terascale computing for scientific research. The laboratories provide state-of-the-art resources for testbeds and novel applications. The division of support between national laboratories and universities is adjusted to maximize scientific productivity.

- ***University Research***: University researchers play a critical role in the nation's research effort and in the training of graduate students. During FY 2004, the ASCR program supported over 150 grants to the nation's university researchers and graduate students engaged in civilian applied mathematics, large-scale network and computer science research. In addition, ASCR supports a Computational Science Graduate Fellowship and an Early Career Principal Investigator activity in Applied Mathematics, Computer Science and High-Performance Networks. In FY 2004, ASCR selected 24 new graduate students representing 17 universities and 13 states, and made 29 awards to early career principal investigators. Approximately half of those who received Ph.D.'s in the Computational Sciences Graduate Fellowship program between 1991 and 2001 are pursuing careers outside universities or national labs. ASCR also provides support to other SC research programs.

The university grants program is proposal driven, similar to the computer science and applied mathematics programs at the National Science Foundation (NSF). However, ASCR grant solicitation notices are focused on topics that have been identified as important for DOE missions. ASCR funds the best among the ideas submitted in response to grant solicitation notices (http://www.sc.doe.gov/grants/grants.html). Proposals are reviewed by external scientific peers and competitively awarded according to the guidelines published in 10 CFR 605 (http://www.science.doe.gov/production/grants/605index.html).

- ***National Laboratory Research***: ASCR supports national laboratory-based research groups at Argonne, Brookhaven, Los Alamos, Lawrence Berkeley, Lawrence Livermore, Oak Ridge, Pacific Northwest, and Sandia national laboratories and Ames Laboratory. The directions of laboratory research programs are driven by the needs of the Department and the unique capabilities of the laboratories to support large scale, multidisciplinary, collaborative research activities. In addition, laboratory-based research groups are highly tailored to the major scientific programs at the individual laboratories and the computational research needs of SC. Laboratory researchers collaborate with other laboratory and academic researchers, and are important for developing and maintaining testbeds and novel applications of high performance computing and networking in SC research. At Los Alamos, Livermore, and Sandia, ASCR funding plays an important role in supporting basic research that can improve the applied programs, such as the Accelerated Strategic Computing Initiative (ASCI) and the Science-Based Stockpile Stewardship program.

ASCR funds field work proposals from the national laboratories. Proposals are reviewed by external scientific peers and awarded using procedures that are equivalent to the 10 CFR 605 guidelines used for the grants program. Performance of the laboratory groups is reviewed by ASCR staff annually to examine the quality of their research and identify needed changes, corrective actions or redirection of effort. Individual laboratory groups have special capabilities or access to laboratory resources that can be profitably utilized in the development of the research program.

Significant Program Shifts

The ASCR program advances mathematics and computer science, and develops the specialized algorithms, the scientific software tools, and the software libraries needed by DOE researchers to effectively use high-performance computing and networking hardware for scientific discovery. The ASCR program has been a leader in the computational sciences for several decades and has been acknowledged for pioneering accomplishments. The FY 2006 ASCR budget is focused in priority areas identified by the Joint OMB-Office of Science and Technology Policy (OSTP) Research Priorities memorandum issued in August 2004.

Major elements of the ASCR portfolio related to SciDAC will be recompeted in FY 2006, with attention paid to support for the long term maintenance and support of software tools such as mathematical libraries, adaptive mesh refinement software, and scientific data management tools developed in the first 5 years of the effort. In addition, in FY 2006 ASCR is changing the way in which it manages its Genomics: GTL partnership with BER. The management of these efforts will be integrated into the portfolio of successful SciDAC partnerships. Finally, in FY 2006 ASCR will initiate a small number of competitively selected SciDAC institutes at universities which can become centers of excellence in high end computational science in areas that are critical to DOE missions at a total funding level of $8,000,000.

The research effort in Collaboratory Tools and Pilots and Networking will be restructured into an integrated Distributed Network Environment activity focused on basic research in computer networks and the middleware needed to make these networks tools for science. This change will enable the reduced NGA effort to operate computers, such as the 20 teraflop Cray X1e and Cray Red Storm system acquired in FY 2004 and FY 2005 at the ORNL-CCS, as tools for science and especially to satisfy the demand for resources that has resulted from the successful SciDAC efforts. In addition, the NGA activity will initiate a new competition for Research and Evaluation Prototype Computer testbeds to enable SciDAC teams to evaluate the potential of future architectures. NGA will continue its focus on research in operating systems and systems software.

The National Leadership Computing Facility acquired under the NGA Leadership Class Competition in FY 2004, will be operated to provide high performance production capability to selected SC researchers. The NGA efforts, as well as the enhancement of NERSC are aligned with the plan developed by the HECRTF established by the National Science and Technology Council (NSTC) and OSTP. These efforts will play a critical role in enabling Leadership Class Machines that could lead to solutions for scientific and industrial problems beyond what would be attainable through a continued simple extrapolation of current computational capabilities. This area has been identified as a priority within the overall Networking and Information Technology Research and Development (NITR&D) priorities of the Administration.

The FY 2006 budget request includes $7,500,000 for continued support of the Genomics: GTL research program, in partnership with the BER program; $2,600,000 for the Nanoscale Science, Engineering and Technology initiative led by the BES program; and $1,350,000 for support of the Fusion Simulation Project, led by the Fusion Energy Sciences program. ASCR's contributions to these partnerships will consist of advancing the mathematics and developing new mathematical algorithms to simulate biological systems and physical systems at the nanoscale. These partnerships support the Biology of Complex Systems and National Nanotechnology Initiative priorities identified by OSTP and OMB.

The FY 2006 budget also includes $8,500,000 to continue the "Atomic to Macroscopic Mathematics" (AMM) research support in applied mathematics needed to break through the current barriers in our understanding of complex physics processes that occur on a wide range of interacting length- and time-scales. Achieving this basic mathematical understanding will provide enabling technology to virtually every challenging computational problem faced by SC.

In FY 2006, the MICS subprogram will continue to support core research activities in applied mathematics, computer science, and network research.

These changes were made to guarantee the continued quality, relevance, and performance of ASCR programs. All ASCR activities undergo prospective and retrospective merit reviews and our extensive use of partnerships with other SC programs ensures the relevance of our efforts to SC missions.

The Laboratory Technology Research subprogram was brought to a successful conclusion in FY 2004 as planned with orderly completion of all existing Cooperative Research and Development Agreements (CRADAs). This does not mean that technology transfer activities have ended; rather, these activities are now institutionalized as a part of the process of doing research at DOE sites.

Interagency Environment

The activities funded by the MICS subprogram are coordinated with other Federal efforts through the Interagency Principals Group, chaired by the President's Science Advisor, and the Information Technology Working Group (ITWG). The ITWG evolved through an interagency coordination process that began under the 1991 High Performance Computing Act as the High Performance Computing, Communications, and Information Technology (HPCCIT) Committee. DOE has been an active participant in these coordination groups and committees since their inception. The MICS subprogram will continue to coordinate its activities through these mechanisms and will lead the development of new coordinating mechanisms as needs arise. The DOE program solves mission critical problems in scientific computing. In addition, results from the DOE program benefit the Nation's information technology basic research effort. The FY 2006 program positions DOE to make additional contributions to this effort. In the area of high performance computing and computation, ASCR has extensive partnerships with other Federal agencies and the NNSA. Examples include: participating in the program review team for the DARPA High Productivity Computing Systems program; serving on the planning group for the Congressionally mandated DOD plan for high performance computing to serve the national security mission; serving on the OSTP High End Computing Revitalization Task Force; and extensive collaboration with NNSA-Advanced Simulation Computing. In FY 2003, ASCR formalized many of these interactions by developing a Memorandum of Understanding with SC, NNSA, DOD's Under Secretary for Defense Research and Engineering, DARPA, and the National Security Administration to coordinate research, development, testing, and evaluation of high performance computers.

Scientific Discovery through Advanced Computing

The Scientific Discovery through Advanced Computing (SciDAC) program is a set of coordinated investments across all SC mission areas with the goal of achieving breakthrough scientific advances via computer simulation that were impossible using theoretical or laboratory studies alone. The power of computers and networks is increasing exponentially. Advances in high-end computing technology, together with innovative algorithms and software, are being exploited as intrinsic tools for scientific discovery. SciDAC has also pioneered an effective new model of multidisciplinary collaboration among discipline-specific scientists, computer scientists, computational scientists, and mathematicians. The product of this collaborative approach is a new generation of scientific simulation codes that can productively exploit terascale computing and networking resources. The program is bringing computation and simulation to parity with experiment and theory in the scientific research enterprise as demonstrated by major advances in climate modeling and prediction, plasma physics, particle physics, accelerator design, astrophysics, chemically reacting flows, and computational nanoscience.

The research focus of ASCR SciDAC activities includes Integrated Software Infrastructure Centers (ISICs). ISICs are partnerships between DOE national laboratories and universities focused on research, development, and deployment of software to accelerate the development of SciDAC application codes. Progress to date includes significant improvements in performance modeling and analysis capabilities that have led to doubling the performance on 64 processors of the Community Atmosphere Model component of the SciDAC climate modeling activity. The three Mathematics ISICs are bringing a new level of mathematical sophistication to computational problems throughout SC. One of these, the Terascale Optimal Partial Differential Equations (PDE) Simulations (TOPS) Center, is combining the Hyper and Portable Extensible Toolkit for Scientific Computation (PETSC) libraries, together with newly developed algebraic multigrid solvers, to create fast algorithms for a variety of tough and important problems, including biochemical reaction diffusion equations and advection equations for combustion simulation. The Terascale Simulation Tools and Technologies Center is working to develop a framework for coupling different types of grids together in a single application. For example, in a simulation of engine combustion, one might want an unstructured grid for the complex geometry around the valves, but a regular grid in the rest of the cylinder. Finally, the Applied Partial Differential Equations Center is focused on using structured adaptive grids for simulation in a variety of application domains, including ground water flow, combustion chemistry, and magneto-hydrodynamics. Given the difficulty of magneto-hydrodynamic simulation, this center is having a strong impact on the design of new particle accelerators.

In FY 2006, ASCR is reducing SciDAC efforts in collaboratories as it restructures and integrates collaboratory research and network research into an integrated effort to develop the middleware and underlying network capabilities needed for the science of the 21st century. The DOE collaboratory research and pilot projects have developed software that enables scientists across the world to access DOE facilities and work together to address critical scientific questions. This software is a critical part of plans in High Energy Physics to access data from the Large Hadron Collider in Switzerland. However, the funds that the reduction in these efforts makes available are needed to enable ASCR to provide adequate high performance computing resources that are critical for SciDAC teams to achieve their full potential. In addition, in FY 2006 ASCR will initiate a small number of competitively selected SciDAC institutes at Universities which can become centers of excellence in high end computational science in areas that are critical to DOE missions at a total funding level of $8,000,000.

Next Generation Computer Architecture for Science and Industry

The Next Generation Computer Architecture for Science and Industry (NGA) research activity is an integral part of an SC strategy to acquire additional advanced computing capability to support existing users in the near-term and to initiate longer-term research and development on next generation computer architectures. The goal of the NGA is to identify and address major architectural bottlenecks in the performance of existing and planned DOE science applications, such as internal data movement in very large systems. In FY 2006, the NGA effort will support computer science research that supports high performance computers such as operating systems that scale to tens of thousands of processors and techniques for evaluating the potential performance of novel computer architectures, operation of high performance computers at ORNL's Center for Computational Sciences, and a new competitive program to select R&E Prototype computer testbeds at multiple locations. The computers at the CCS will enable significant scientific progress by delivering significant increases in performance to critical DOE mission applications; and will also enable industrial researchers to find opportunities for virtual prototypes and simulation of industrial processes that result in an enhanced competitive position because of sharply reduced 'time to market'.

The new R&E Prototype testbeds will enable SciDAC teams and other leading edge computational scientists to evaluate the potential of new computer architectures as tools for science.

The NGA activity is coordinated with other Federal agencies to gain additional insight into research directions, optimize the utilization of resources, and establish the framework for a national effort. In May 2004, the Administration issued the report of the HECRTF, titled "Federal Plan for High-End Computing," which specifically cites DOE's NGA activity. The Plan emphasizes a coordinated, sustained research program for high-end computing and an interagency collaborative strategy for addressing mission agency needs for additional computational resources. SC has a prominent role throughout this plan, and the NGA effort is fully aligned with the plan.

Scientific Facilities Utilization

The ASCR program's FY 2006 request includes support to the NERSC, ESnet, and Center for Computational Sciences (CCS), components of SC-wide Facilities Optimization effort. The investment in NERSC will provide computer resources for about 2,000 scientists in universities, DOE laboratories, Federal agencies, and U.S. companies. The proposed funding will enable NERSC to maintain its role as one of the Nation's premier unclassified computing centers, a critical element for success of many SC research programs. The investment in ESnet will provide the DOE science community with capabilities not available through commercial networks or the commercial Internet. ESnet provides national and international high-speed access to DOE and SC researchers and research facilities, including light sources, neutron sources, particle accelerators, fusion reactors, spectrometers, supercomputers, and other high impact scientific instruments. The investment in the CCS will provide operation of the Leadership Class Computing Capability for Science.

	FY 2004	FY 2005 Est.	FY 2006 Est.
NERSC			
Maximum Hours	8,760	8,760	8,760
Scheduled Hours	8,585	8,585	8,585
Unscheduled Downtime	1%	1%	1%
ESnet			
Maximum Hours	8,760	8,760	8,760
Scheduled Hours	8,585	8,585	8,585
Unscheduled Downtime	1%	1%	1%
CCS			
Maximum Hours	7,008	7,008	7,008
Scheduled Hours	7,008	7,008	7,008
Unscheduled Downtime	1%	1%	1%

Workforce Development

The R&D Workforce Development mission is to ensure the supply of computational and computer science Ph.D. level scientists for the Department and the Nation through graduate student and postdoctoral research support. In FY 2006, this program will support approximately 800 graduate students and post doctoral investigators, of which 500 will be supported at SC user facilities.

ASCR will continue the Computational Science Graduate Fellowship Program with the successful appointment of 20 new students to support the next generation of leaders in computational science.

	FY 2004	FY 2005 est.	FY 2006 est.
# University Grants	140	142	135
Average Size/Duration	$197,000/yr-3yrs	$197,000/yr-3yrs	$197,000/yr-3yrs
# Laboratory Groups	165	165	155
# Graduate Students (FTEs)	354	354	350
# Permanent Ph.D.s (FTEs)	675	675	625

Mathematical, Information, and Computational Sciences

Funding Schedule by Activity

	(dollars in thousands)				
	FY 2004	FY 2005	FY 2006	$ Change	% Change
Mathematical, Information, and Computational Sciences					
Mathematical, Computational, and Computer Sciences and Distributed Network Environment Research	97,912	102,575	107,026	+4,451	+4.3%
High Performance Computing and Network Facilities and Testbeds	95,967	123,608	94,400	-29,208	-23.6%
SBIR/STTR	0	6,285	5,629	-656	-10.4%
Total, Mathematical, Information, and Computational Sciences	193,879	232,468	207,055	-25,413	-10.9%

Description

The Mathematical, Information, and Computational Sciences (MICS) subprogram is responsible for carrying out the mission of the ASCR program: To deliver forefront computational and networking capabilities to scientists nationwide that enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy.

Benefits

MICS supports ASCR's contribution to DOE's mission to provide world-class scientific research capacity by providing world-class, peer-reviewed scientific results in mathematics, high performance computing and advanced networks and applying the potential of terascale computing to advanced scientific applications. Computer-based simulation enables us to predict the behavior of complex systems that are beyond the reach of our most powerful experimental probes or our most sophisticated theories. Computational modeling has greatly advanced our understanding of fundamental processes of nature, such as fluid flow and turbulence or molecular structure and reactivity. Through modeling and simulation, we will be able to explore the interior of stars and learn how protein machines work inside living cells. We can design novel catalysts and high-efficiency engines. Computational science is increasingly central to progress at the frontiers of almost every scientific discipline and to our most challenging feats of engineering. The science of the future demands that we advance beyond our current computational abilities. Accordingly, we must address the following questions:

- What new mathematics are required to effectively model systems such as the earth's climate or the behavior of living cells that involve processes taking place on vastly different time and/or length scales?
- Which computational architectures and platforms will deliver the most benefit for the science of today and the science of the future?
- What advances in computer science and algorithms are needed to increase the efficiency with which supercomputers solve problems for SC?

- What operating systems, data management, analysis, model development, and other tools are required to make effective use of future-generation supercomputers?
- Is it possible to overcome the geographical distances that often hinder science by making all scientific resources readily available to scientists, regardless of whether they are at a university, national laboratory, or industrial setting?

To answer these questions and develop the algorithms software and tools that are needed, MICS has developed four strategies (the strategy numbers refer to the SC Strategic Plan):

6.1 Advance scientific discovery through research in the computer science and applied mathematics required to enable prediction and understanding of complex systems.

6.2 Extend the frontiers of scientific simulation through a new generation of computational models that fully exploit the power of advanced computers and collaboratory software that makes scientific resources available to scientists anywhere, anytime.

6.3 Bring dramatic advances to scientific computing challenges by supporting the development, evaluation, and application of supercomputing architectures tailored to science.

6.4 Provide computing resources at the petascale and beyond, network infrastructure, and tools to enable computational science and scientific collaboration.

All MICS investments directly contribute to one or more of these strategies.

Supporting Information

The computing and the networking capabilities required to meet SC needs exceed the state-of-the-art by a wide margin. Discussions of the extent to which the requirements of science exceed the current state-of –the-art can be found in a number of reports including: “Federal Plan for High-End Computing, Report of the High-End Computing Revitalization Task Force (HECRTF)” May 10, 2004, Appendices A-1, A-2, and A-3, (http://www.sc.doe.gov/ascr/hecrtfrpt.pdf); “A Science-Based Case for Large-Scale Simulation,” Volume 1, July 30, 2003, (http://www.sc.doe.gov/ascr/Scalesreptvol1.pdf); “Theory and Modeling in Nanoscience, Report of the May 10-11, 2002 Workshop conducted by the Basic Energy Sciences and Advanced Scientific Computing Advisory Committees to the Office of Science, Dept. of Energy,” (http://www.sc.doe.gov/ascr/TMN_rpt.pdf); “Integrated Simulation and Optimization of Magnetic Fusion Systems, Report of the FESAC Panel,” November 2, 2002. (http://www.ofes.fusion.doe.gov/News/FSP_report_Dec9.pdf); and “High-Performance Networks for High-Impact Science, Report of the High-Performance Network Planning Workshop conducted August 13-15, 2002,” (http://www.sc.doe.gov/ascr/high-performance_networks.pdf). Furthermore, the algorithms and software tools, libraries, and environments needed to accelerate scientific discovery through modeling and simulation are beyond the realm of commercial interest. To establish and maintain DOE’s modeling and simulation leadership in scientific areas that are important to its mission, the MICS subprogram employs a broad, but integrated, research strategy. The MICS subprogram’s basic research portfolio in applied mathematics and computer science provides the foundation for enabling research activities, which includes efforts to advance networking and to develop software tools, libraries, and environments. Results from enabling research supported by the MICS subprogram are used by computational scientists supported by other SC and DOE programs. This link to other DOE programs provides a tangible assessment of the value of the MICS subprogram for advancing scientific discovery and technology development through simulations.

In addition to its research activities, the MICS subprogram plans, develops, and operates supercomputer and network facilities that are available—24 hours a day, 365 days a year—to researchers working on problems relevant to DOE's scientific missions.

The Early Career Principal Investigator (ECPI) activity was initiated in FY 2002 for scientists and engineers in tenure track positions at U.S. universities. Seventeen awards were made in FY 2002, thirteen awards in FY 2003, and sixteen in FY 2004. Additional awards will be made in FY 2005 for this activity, pending peer review of applications. The goal of the ECPI activity is to support SC mission-related research in applied mathematics, computer science, and high-performance networks performed by exceptionally talented university investigators, who are at an early stage in their professional careers.

FY 2004 Accomplishments

- A collaboratory pilot project, the Collaboratory for Multi-scale Chemical Sciences (CMCS), has ushered in a new era of validated chemical reference data where all the pertinent experiments and computations can be considered by all the experts. A fifty-year-old question of the vaporization enthalpy of graphite, a thermochemical reference value for countless computational and experimental studies, has been resolved by an International Union of Pure and Applied Chemistry (IUPAC) Task Group empowered by CMCS capabilities. The new value for the enthalpy of formation of carbon atom in the gas phase is now more than twice as accurate as the previously accepted value. A new CMCS application called Active Thermochemical Tables enabled IUPAC group members to collaborate on the systematic reevaluation of previous experimental data, new state-of-the-art computational chemistry results, and comparison with other data. These new capabilities also enabled a group at DOE's Argonne National Laboratory to fully confirm a recent revision of the enthalpy of formation of the pivotal combustion and atmospheric radical, hydroxyl (OH), and to further reduce its uncertainty by a factor of ~6.5, thus removing a potential source of uncertainty in current chemical models.
- Under the Particle Physics Data Grid (PPDG) collaboratory project, robust, sustained, hands-off, production data transfer of terabytes of data has been enabled using GridFTP and Storage Resource Manager (SRM) implementations. Scientists on PPDG experiments moved terabytes (TB) of data routinely between institutions in the collaboration at a higher data transfer rate, a higher success rate for the transfer, high reliability in the resulting cataloging, and a decreased manpower effort in achieving the data transfer. This enables faster turnaround for analysis of the data, earlier physics results, and decreased risk in the building of new detectors and computing systems. The STAR detector at Brookhaven National Laboratory's (BNL's) Relativistic Heavy Ion Collider (RHIC) has seen a factor of ten increase in throughput for moving files between BNL and LBNL. A sustained robust, automated 5 TB data movement per week allows "next day" analysis. A factor of fifty improvement has been made in file discrepancies (now < 0.02% error rate). Publication of first results of the d+Au run last spring appeared in print in a record four months from the end of the run. The DØ experiment at the Fermi National Accelerator Laboratory (FNAL) moved over 50 TB of event data (20% of the run) to be analyzed off-site over the past six months. By using multiple streams in GridFTP, a factor of five improvement in throughput was made possible.
- The Earth System Grid (ESG) collaboratory project now provides the climate research community with a powerful new capability, and an excellent Grid-based foundation upon which to rapidly expand scientific services. The ESG has developed new data management capabilities that provide robust interoperability among DOE (High Performance Storage System [HPSS]) and National

Center for Atmospheric Research (NCAR) (Mass Storage System [MSS]) archival systems. It is now possible for a climate researcher to securely access through the ESG web portal all the climate modeling datasets (about 100) that have been developed over the past five years and download the portion of the dataset of interest. Additionally, ESG has been chosen as the vehicle for delivering climate model data to the international community for the fourth Intergovernmental Panel on Climate Change (IPCC) Assessment. The Program for Climate Model Diagnosis and Intercomparison (PCMDI), at the request of the World Meteorological Organization (WMO), hosted the primary ESG/IPCC services in the summer timeframe, 2004.

- Language Interoperability Improvements from Common Component Architecture — The Babel interoperability language allows software to be written in a variety of computer languages. Any one of C, C++, and Python can be used interchangeably in the same program. Recently Fortran, a language particularly important to scientific computing has also been added as a full peer to the rest of the Babel languages. Now software libraries using Babel are automatically callable from Fortran simulation codes. Conversely, Fortran teams now have a new capability to export their software to applications written in other languages. For example, all of the Common Component Architecture (CCA) frameworks and CCA-compliant component libraries are not only callable from Fortran, but the Fortran community can participate in developing fully CCA compliant components of their own. A concrete example of how Babel's language technology impacted science can be found in CCA's molecular geometry optimization activities. North West Chem (NWChem), Massively Parallel Quantum Chemistry (MPQC), Portable Extensible Toolkit for Scientific Computation (PETSC) and The Adaptive communication environment Object request broker (TAO) were combined into a single application with demonstrated improvements in robustness and efficiency over classic techniques. Babel was used in this effort since NWChem is written in Fortran, and MPQC in C++, while PETSC and TAO are written in C.

- Improved performance modeling framework predicted performance within 15% for scientific application codes—A performance model is an encapsulation of the performance characteristics of a given scientific application on a given computer system, which then can be used to predict the performance of the code on a new system. An accurate and easy-to-use performance modeling capability has many potential applications, including more effective system design, streamlined system procurements, and even easier tuning of programs by applications scientists. As part of the Performance Evaluation Research Center (PERC) SciDAC ISIC, University of California/San Diego researchers have developed a methodology for performance modeling that significantly simplifies this process, replacing in-depth analysis by automated tools. They developed and applied a new, faster and more accurate MetaSim framework to several application codes, and achieved performance predictions accurate to within 15%. In particular, the predictions for Hybrid Coordinate Ocean Model (HYCOM) and Cobalt60 were performed at various problem sizes for processor counts ranging from 64 to 256 on three different architectures.

- Rose Code Optimization—The Rose source code optimization project at Lawrence Livermore National Laboratory (LLNL) now includes support for all C++ language constructs, C, and United Parallel Code (UPC), and work to support Fortran 90 has begun in collaboration with Rice University. Performance increases up to a factor of 6 have been obtained by optimizing high level abstractions. Additional increases on large array operations, up to a factor of 15, have been obtained by combining high-level abstractions with aggressive loop optimizations. Program analysis of projects with more than a million lines of code and hundreds of files has been enabled through extensive interactive visualization techniques and new database features that support storage of intermediate analysis results. This work impacts national and international research programs through strong collaborations with 14 universities and 4 DOE laboratories.

- High Performance Scientific Software Components—Argonne computer scientists, in close collaboration with the SciDAC Center for Component Technology for Terascale Simulation Software and the Common Component Architecture ISIC, have developed a variety of high-performance numerical components that are being used in prototype Partial Differential Equation (PDE)-based simulations as well as in full-blown applications in computational chemistry and climate modeling. Prototype software enables the automatic selection and configuration of components to suit the computational conditions imposed by a simulation and its operating environment. In collaboration with chemists at PNNL and Sandia National Laboratory (SNL), optimization components have been employed in the fundamental problem of molecular geometry optimization. Components based on the Toolkit for Advanced Optimization provide robust and efficient optimization solvers, which in turn can employ the parallel linear algebra capabilities within the PETSC. Experiments on a representative set of chemical structures have demonstrated improved robustness and efficiency.
- Design Patterns for Parallel Programming—Design patterns are an important tool for organizing and simplifying complex code development. The University of Illinois at Urbana-Champaign is identfying key programming patterns that occur in parallel scientific application, and to date 10 patterns have been documented: Pipes, Layers, Repositories, Master-Worker, Replicable, Inseparable, Divide and Conquer, Geometric, Recursive, and Irregular Mesh. Each of these patterns has been analyzed in terms of the decomposition paradigm (data decomposition vs. functional decomposition), ordering constraints, communication structure (local, tree, and irregular), and dependency structure (separable, deterministic/function dependency, and inseparable); and each of them has been documented in the style used by programming patterns dictionaries.
- Open-source compiler for Co-array Fortran—Co-Array Fortran (CAF) is a model for parallel programming that consists of a small set of extensions to the Fortran 90 programming language, which is widely used within the DOE community for scientific computing. CAF shifts the burden for managing the details of communication between a computer system's processors from the application developer to the CAF compiler. In January 2004, Rice University released an open-source CAF compiler prototype for evaluation by the science community and vendors. This project has already attracted significant interest from vendors of high performance computing systems. An ongoing effort is focused on refining the CAF language to further simplify writing scalable parallel programs and devising more sophisticated compiler technology to deliver highest possible performance for scientific computations on parallel systems.
- Scientific Application Performance on Modern Parallel Vector Systems—Researchers at LBNL are conducting a study to evaluate four diverse scientific applications from plasma physics, material sciences, astrophysics, and magnetic fusion. Performance comparisons have been made between vector-based (Earth Simulator and Cray X1), and leading superscalar-based (IBM Power3/4 and SGI Altix) platforms. This is the first international group to conduct a performance evaluation study at the Earth Simulator Center. Results demonstrate that the vector systems achieve excellent performance on the application suite—the highest of any architecture tested to date.

- Wide acceptance of open-source, high-end cluster software by industry and users—The ORNL Open Source Cluster Application Resources (OSCAR) cluster computing software for high-end computing continues to expand its capability and to increase its user base. The software has been downloaded by more than 130,000 groups around the world and is promoted by cluster vendors including Dell and Intel. The adoption of this system has expanded the number of software packages available to the cluster community, and OSCAR continues to reduce cluster total cost of ownership. It has simplified the job of software authors, system administrators, and ultimately the application user by providing a timely and much simpler method of supplying and applying software updates. The SciDAC Scalable Systems Software ISIC leverages OSCAR technology to simplify deployment for the end-user as well as application developers.

FY 2004 Awards

- R&D 100 Award to Los Alamos National Laboratory (LANL) Team — Every year, R&D Magazine recognizes the world's top 100 scientific and technological advances with awards for innovations showing the most significant commercial potential. The Computer Science project "Science Appliance" received an R&D 100 award for the year 2004. The Science Appliance software suite, called Clustermatic, is a set of tools that revolutionizes the way clustered systems are managed, monitored, administered, and run. Clustermatic increases reliability and efficiency, decreases node autonomy, and reduces administration costs, enabling commodity-based cluster systems to be more competitive with high-end supercomputers. Clustermatic also won the "Best Open-Source Cluster Solution" award at Cluster World 2004.
- Researchers at ANL and Northwestern University were awarded the Beale-Orchard-Hays Prize for Excellence in Computational Mathematical Programming. This prestigious prize is bestowed once every three years and is considered a major award from the Mathematical Programming Society. The presentation took place at the 18th International Symposium on Mathematical Programming in Copenhagen. The Argonne team received the award in recognition of their fundamental research in support of the development of the NEOS (Network-Enabled Optimization System) Server. Today's NEOS Server is a collaborative project that provides access to dozens of academic and commercial optimization solvers through an assortment of Internet interfaces. Over 120,000 job requests are handled annually, including optimization problems from academic, commercial, and government institutions.
- Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program—Some of the most significant work done at NERSC in 2004 was made possible by the INCITE Program, which supports a small number of computationally intensive large-scale research projects that are expected to make high-impact scientific advances through the use of a substantial allocation of computer time and data storage at the NERSC Center. In December 2003, SC selected three computational science projects to receive a total of 4.9 million hours of supercomputing time at the NERSC Center—10% of the total computing time available in FY 2004 on NERSC's Seaborg system. One of the INCITE projects, "Thermonuclear Supernovae: Stellar Explosions in Three Dimensions," has achieved unprecedented simulations of supernovae using NERSC's computing resources. The research group at the Center for Astrophysical Thermonuclear Flashes at the University of Chicago, typically uses 512 to 1,024 processors to test their applications. They expect to use 4,096 or more processors for their final calculations. So far they have calculated three Type Ia supernova explosion models, including one octant model with 8 kilometer resolution, which is

considered state of the art in this field. The group has also calculated one full-star model with 8 kilometer resolution and 30 kilometer ignition regions and one full-star model with 8 kilometer resolution and 50 kilometer ignition regions. These last two runs are the first of this kind ever calculated—no other group has produced full-star simulations before. They are now moving to 4 kilometer and 1 kilometer resolution models.

Detailed Justification

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
Mathematical, Computational, and Computer Sciences and Network Environment Research	**97,912**	**102,575**	**107,026**
▪ **Applied Mathematics**..	**22,553**	**26,428**	**28,995**

This activity supports research on the underlying mathematical understanding of physical, chemical, and biological systems, and on advanced numerical algorithms that enable effective description, modeling, and simulation of such systems on high-end computing systems. It directly supports SC Strategic Plan strategy 6.1. Research in Applied Mathematics supported by the MICS subprogram underpins computational science throughout DOE. Historically, the numerical algorithms developed under this activity have produced scientific advances through simulation that are as significant as those resulting from improvements in computer hardware. This activity supports research at DOE laboratories, universities, and private companies. Many of the projects supported by this activity involve research partnerships between DOE's national laboratories and universities. The activity supports research in a wide variety of areas of mathematics, including: ordinary and partial differential equations and solution methods, including techniques to convert equations into discrete elements and boundary integral methods; advanced treatment of interfaces and boundaries (fast marching and level set methods, and front tracking); numerical linear algebra (advanced iterative methods, general and problem-specific preconditioners, sparse solvers, and dense solvers); fluid dynamics (compressible, incompressible, and reacting flows, turbulence modeling, and multiphase flows); optimization (linear and nonlinear programming, interior-point methods, and discrete and integer programming); mathematical physics; control theory (differential-algebraic systems, order reduction, and queuing theory); accurate treatment of shock waves; "fast" methods (fast multipole and fast wavelet transforms); mixed elliptic-hyperbolic systems; dynamical systems (chaos theory, optimal control theory, and bifurcation theory); automated reasoning systems; and multiscale mathematics. This final area represents our most recent effort at focusing on those mission-related applications which span wide ranges of interacting length- and time-scales.

The FY 2006 budget continues the Computational Sciences Graduate Fellowship program at the current level of $3,500,000. The FY 2006 budget also includes $8,500,000 for the Atomic to Macroscopic Mathematics effort, an increase of $2,567,000 over the level planned for FY 2005. This increase will support a full year of funding for projects initiated in FY 2005.

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
▪ **Computer Science**	**19,402**	**19,909**	**24,380**

This activity supports research in computer science to enable computational scientists to effectively utilize high-performance computers to advance science in areas important to the DOE mission. This activity supports SC Strategic Plan strategies 6.1 and 6.3. DOE has unique requirements for high performance computing that significantly exceed the capability of software products from computer vendors. This activity supports computer science research in two general areas: the underlying software to enable applications to make effective use of computers with hundreds or thousands of processors as well as computers that are located at different sites; and large scale data management and visualization for both local data analysis and for circumstances where key resources and users are geographically distributed. Research areas include: tools to monitor the performance of scientific applications and enable users to improve performance and get scientific results faster; new programming models to simplify application code development; advanced techniques for visualizing very large-scale scientific data; and Next Generation Architecture (NGA) efforts to improve application performance through innovative next generation operating systems. Researchers at DOE laboratories and universities, often working in partnerships, propose and conduct this research.

All of the computer science research funded by this activity is reported to and coordinated through the High End Computing Coordinating Group of the Interagency Working Group on Information Technology Research and Development. The quality, relevance and performance of the program is continually monitored through extensive peer review; interagency reporting and coordination; and interaction with end users to assist in the determination of impact and future research priorities.

Beginning in FY 2004, this activity incorporated the software research component of NGA to improve application performance and system reliability through innovative approaches to next generation operating systems. In FY 2006, the total funding for the NGA software research component is $6,659,000. This activity is coordinated with other agencies through the High End Computing University Research Activity (HEC URA), an outgrowth of the HECRTF. These activities will be modestly increased in FY 2006, especially in areas such as: performance analysis of innovative high-end architectures; frameworks for data intensive and visual computing; intelligent program development environments; application-specific problem solving environments; and common compile and runtime infrastructures and interfaces, where ASCR is the leader in the Federal agency research efforts. This research will play a key role in the interagency strategy for high end software development recommended in the HECRTF plan.

▪ **Computational Partnerships**..................................	**33,300**	**34,876**	**39,887**

This activity supports the amalgam of those activities previously titled "Advanced Computing Software Tools" and "Scientific Applications Partnerships." The advanced computing software tools part of this activity supports research and development activities that extend key results from applied mathematics and computer science research to develop integrated software tools that computational scientists can use in high performance scientific applications (such as characterizing and predicting phase changes in materials). These tools, which enable improved performance on high-end systems, are critical to the ability of scientists to attack the complex scientific and engineering problems that can only be solved with such systems. This activity directly supports SC Strategic Plan strategy 6.2.

In FY 2006, this activity will support the completion of the original Integrated Software Infrastructure Centers (ISICs) SciDAC activity, competitively selected in FY 2001. The ISICs funded under this activity focus on several important computational infrastructure problems including: structured and unstructured mesh generation for large simulations and high performance tools for solving partial differential equations on parallel computers; tools for analyzing the performance of scientific simulation software that uses thousands of processors; the development of data management and visualization software capable of handling terabyte scale data sets extracted from petabyte scale data archives; software for managing computers with thousands of processors; and software component technology to enable rapid development of efficient, portable, high performance parallel simulation software.

The ISICs are a fundamental component in DOE's SciDAC strategy. The ISICs are responsible for the entire lifecycle of the software that they develop. These software tools must be reliable, understandable, and well documented. Also, the scientific user community needs these tools to be maintained, bug-free, and upgraded as necessary. Since software tools for high performance scientific simulations have no commercial market, the ISICs provide the only means for developing and deploying these tools to the scientific community.

The scientific applications partnerships part of this activity, formerly titled Scientific Application Pilot Projects, supports collaborative research with computational scientists in other disciplines to apply the computational techniques and tools developed by other MICS activities to address problems relevant to the SC mission. This effort tests the usefulness of advances in computing research, transfers the results of this research to the scientific disciplines, and helps define opportunities for future research. The FY 2006 funding for this activity will allow the continuation of the multidisciplinary partnerships that were competitively selected in FY 2001. These projects are part of the SciDAC activity and are coupled to the ISICs. Areas under investigation include design of particle accelerators with the High Energy Physics (HEP) and Nuclear Physics (NP) programs, plasma turbulence in tokamaks with the Fusion Energy Sciences (FES) program, global climate change with the Biological and Environmental Research (BER) program, and combustion chemistry with the Basic Energy Sciences (BES) program. This activity directly supports SC Strategic Plan strategy 6.2.

The FY 2006 request includes funds to continue the partnerships with the BER Genomics: GTL program, the BES program in nanoscale science, and the FES program for the Fusion Simulation Project (FSP).

The FY 2006 request also includes $8,000,000 to initiate two competitively selected SciDAC institutes at Universities which can become centers of excellence in high end computational science in areas that are critical to DOE missions.

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
▪ **Distributed Network Environment Research**	**22,657**	**21,362**	**13,764**

This activity supports the integration of activities previously described under the titles: "Network Research," "Collaboratory Tools," and "Collaboratory Pilots." This integrated activity builds on results of fundamental research in computer science and networking to develop an integrated set of software tools and services to support distributed scientific collaborations and provide end-to-end network performance well beyond the levels that can be achieved today. It advances the Network Environment vision by supporting research and development for advanced capabilities and technologies. For example, it includes standards-based network protocols and middleware that address challenging issues such as security, information location, and network performance that are encountered with ultra-high-speed data transfers, remote visualization, real-time remote instrumentation, and large-scale scientific collaboration. These tools provide a new way of organizing and performing scientific work, e.g., distributed teams and real-time remote access to SC facilities that offers the potential for increased productivity and efficiency. It will also enable broader access to important DOE facilities and data resources by scientists and educators throughout the country. It is particularly important to provide for efficient, high-performance, reliable, secure, and policy-aware management of large-scale data movement across the research enterprise.

Although this activity directly supports SC Strategic Plan strategy 6.2, it is being reduced in FY 2006 to enable MICS to provide computing resources to scientists across SC.

High Performance Computing and Network Facilities and Testbeds	**95,967**	**123,608**	**94,400**
▪ **High Performance Computing Facilities and Testbeds**	**79,387**	**105,320**	**76,112**

This activity represents the integration of activities previously described separately as NERSC and ACRTs. It includes NERSC ($37,868,000), resources at the ORNL Center for Computational Sciences (CCS) ($25,000,000), as well as a new activity to support Research and Evaluation Prototypes that enable SciDAC teams to evaluate the promise of future computer architectures for their applications ($13,244,000). These new testbeds will be selected competitively. This activity directly supports SC Strategic Plan strategy 6.4 through a portfolio of capabilities that range from Research and Evaluation Prototypes (R&E Prototypes) to Leadership Class Computers (LCCs) to High Performance Production Computing (HPPC). HPPC includes NERSC as well as resources at the ORNL CCS, which have completed their evaluation as R&E prototypes but have significant capability that can be cost-effectively delivered to science. This restructured approach will allow for a comparable level of high performance production computing capability to be provided more efficiently.

	FY 2004	FY 2005	FY 2006
	(dollars in thousands)		
• **National Energy Research Scientific Computing Center (NERSC)**	**32,906**	**37,868**	**37,868**

NERSC, located at the LBNL, delivers high-end capability computing services and support to the entire DOE SC research community. NERSC provides these services to the DOE community, to the other DOE laboratories, and to major universities performing work relevant to DOE missions. NERSC provides the majority of resources and services that are used to support SC SciDAC programs. The center serves 2,000 users working on about 700 projects. 35% of users are university based, 61% are in National Laboratories, 3% are in industry, and 1% are in other government laboratories. The major computational resource at NERSC is an IBM SP computer called NERSC 3e. The FY 2006 funding will support the continued operation of NERSC 3e at 10 teraflops peak performance, and a new computer system, which is focused on high performance capacity for scientific applications that do not scale well to more than 512 processors and are not well suited to the IBM SP with a peak performance over 6 teraflops. In addition, in FY 2006 a procurement will be conducted for the next generation of high performance resources at NERSC to be delivered in early FY 2007. These computational resources are integrated by a common high performance file storage system that enables users to easily use all the resources. The FY 2006 budget at NERSC will result in the elimination of some research activities, such as advanced computer architectures, numerical algorithms, and high performance file systems, which had previously been included in the NERSC budget. These research activities will now be completed within the peer reviewed computer science, applied mathematics, and partnership budget activities described earlier. This NERSC budget will maintain NERSC's world leading level of user support and enable introduction of new computing capacity to support critical science needs.

• **Center for Computational Sciences (CCS)**	**46,481**	**67,452**	**25,000**

The CCS, located at ORNL, provides high-end capability computing services to SciDAC teams and other DOE users on Eagle, a 720 processor IBM SP3 system and Cheetah, an 864 processor IBM SP4 system. As in the case at NERSC, these computational resources are integrated by a common high performance file storage system that enables users to easily use all the computational resources.

In FY 2004, the DOE leadership-class computing capability for science activity was initiated at CCS as the result of an open competition, and we anticipate that these facilities will become a major capability for science in the late FY 2005 timeframe with a Cray X1E system, the most capable system available to scientific users in the U.S., and a complementary Cray Red Storm System. These computers will become a part of the overall SC portfolio of HPPC resources in FY 2006. In FY 2006, the CCS will operate the computers acquired in FY 2004 and FY 2005 as resources for science. In order to enable effective operation of these resources for SciDAC teams and other scientists, all FY 2005 activities focused on the acquisition, testing, and evaluation of R&E prototype computer hardware testbeds to assess the prospects for meeting future computational needs of SC will be concluded by the end of FY 2005. Possible future

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006

R&E prototype activities at the CCS will be competitively selected within the new Research and Evaluation Partnerships activity described below.

Related requirements for capital equipment, such as high-speed disk storage systems, archival data storage systems, and high performance visualization hardware, and general plant projects (GPP) funding are also supported. FY 2006 capital equipment requirements for these types of capital equipment will be reduced from FY 2005.

• **Research and Evaluation Prototypes**	**0**	**0**	**13,244**

This new activity in FY 2006 will support Research and Evaluation prototypes to enable SciDAC teams to evaluate the potential of new computer architectures as tools for Science. These new testbeds will be competitively selected.

▪ **High Performance Network Facilities and Testbeds**...	**16,580**	**18,288**	**18,288**

This activity supports the ESnet, a Wide Area Network (WAN) project that supports the scientific research mission of the DOE. The ESnet project/investment supports the agency's mission and strategic goals and objectives by providing DOE with interoperable, effective, and reliable communications infrastructure and leading-edge network services. ESnet supplies the DOE science community with capabilities not available through commercial networks or the commercial Internet. ESnet provides national and international high-speed access to DOE and SC researchers and research facilities, including light sources, neutron sources, particle accelerators, fusion reactors, spectrometers, supercomputers, and other high impact scientific instruments. ESnet provides the communications fabric that interconnects geographically distributed research facilities and large-scale scientific collaborations. ESnet supplies the critical infrastructure that links DOE researchers worldwide and forms the basis for advanced experimental research in networking, collaboratory tools, and distributed data-intensive scientific applications testbeds such as the national collaboratory pilot projects. This activity directly supports SC Strategic Plan strategy 6.4. ESnet provides network services through contracts with commercial vendors for advanced communications services including Asynchronous Transfer Mode (ATM), Synchronous Optical Networks (SONET) and Dense Wave Division Multiplexing (DWDM). ESnet provides interfaces between the network fabric it provides and peering arrangements with other Federal, education, and commercial networks, international research network connections, and the University Corporation for Advanced Internet Development (UCAID) Abilene network that provides high performance connections to many research universities.

In FY 2006, funds will be used to operate ESnet and to continue support for upgrading the capability of the ESnet backbone to over 10 Gigabits per second (Gbps) from its current capability of 1 Gbps. ESnet will control costs for connections of ESnet sites to the ESnet backbone by partnering with other organizations to develop metropolitan area fiber networks. Remaining funds will be used to upgrade networking hardware and services at high priority ESnet sites to exploit the enhanced performance capabilities of the backbone. FY 2006 capital equipment requirements remain at the same level as in FY 2005.

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
Small Business Innovative Research (SBIR)/ Small Business Technology Transfer (STTR)	**0**	**6,285**	**5,629**

In FY 2004 $4,833,000 and $580,000 were transferred to the SBIR and STTR programs respectively. The FY 2005 and FY 2006 amounts shown are the estimated requirements for the continuation of the SBIR and STTR program.

	FY 2004	FY 2005	FY 2006
Total, Mathematical, Information, and Computational Sciences	**193,879**	**232,468**	**207,055**

Explanation of Funding Changes

	FY 2006 vs. FY 2005 ($000)
Mathematical, Computational & Computer Sciences and Distributed Network Environment Research	
▪ **Applied Mathematics**	
Increase in Applied Mathematics activity to support full year of funding for projects initiated in FY 2005 as part of Atomic to Macroscopic Mathematics effort.	+2,567
▪ **Computer Science**	
Enhance NGA computer science research to enable effective software for high performance computers.	+4,471
▪ **Computational Partnerships**	
Increase in partnership activities resulting from recompetition of SciDAC activities and initiation of new university based competition for SciDAC Institutes which can become centers of excellence in scientific areas critical to the missions of DOE.	+5,011
▪ **Distributed Network Environment Research**	
This activity is being reduced and focused on network research and middleware to enable MICS to provide computer resources to scientists across the SC.	-7,598
Total, Mathematical, Computational & Computer Sciences and Distributed Network Environment Research	**+4,451**

	FY 2006 vs. FY 2005 ($000)
High Performance Computing & Network Facilities and Testbeds	
▪ **High Performance Computing Facilities and Testbeds**	
This reduction reflects the completion in FY 2005 of evaluation activities at the CCS, rescoping of the Leadership Class effort to operate computers acquired in prior years as resources for science, and initiation of a new competitive program to select promising future architecture computing testbeds.	-29,208
SBIR/STTR	
Decrease in SBIR/STTR due to decrease in operating expenses.	-656
Total Funding Change, Mathematical, Information, and Computational Sciences	**-25,413**

Laboratory Technology Research

Funding Schedule by Activity

(dollars in thousands)

	FY 2004	FY 2005	FY 2006	$ Change	% Change
Laboratory Technology Research					
Laboratory Technology Research	2,916	0	0	0	0.0%
SBIR/STTR	0	0	0	0	0.0%
Total, Laboratory Technology Research	2,916	0	0	0	0.0%

Description

The Laboratory Technology Research (LTR) subprogram is brought to a successful conclusion in FY 2004 with orderly completion of all existing Cooperative Research and Development Agreements (CRADAs). The mission of the LTR subprogram was to support high-risk research that advances science and technology to enable applications that could significantly impact the Nation's energy economy. LTR fostered the production of research results motivated by a practical energy payoff through cost-shared collaborations between SC laboratories and industry. The termination of the LTR subprogram does not mean that technology transfer activities have ended; rather, due to the impact of this subprogram, these activities are now institutionalized as a part of the process of doing research at DOE sites.

Benefits

LTR supported ASCR's contribution to DOE's mission of world-class scientific research capacity by promoting the transfer of these research results to the private sector. The success of this program has institutionalized these processes in all of the programs within SC; therefore, these processes are now integrated into the other programs and the LTR subprogram is no longer needed.

Detailed Justification

(dollars in thousands)

	FY 2004	FY 2005	FY 2006
Laboratory Technology Research	**2,916**	**0**	**0**

This activity supported research to advance the fundamental science at SC laboratories toward innovative energy applications. Through CRADAs, the SC laboratories entered into cost-shared research partnerships with industry, typically for a period of three years, to explore energy applications of research advances in areas of mission relevance to both parties. The existence of the LTR subprogram fostered the institutionalization of technology transfer activities at DOE sites. Now that these activities are institutionalized, a separate program to fund them is no longer necessary.

	(dollars in thousands)		
	FY 2004	FY 2005	FY 2006
SBIR/STTR	**0**	**0**	**0**

In FY 2004, $75,000 and $9,000 were transferred to the SBIR and STTR programs respectively.

Total, Laboratory Technology Research	**2,916**	**0**	**0**

Capital Operating Expenses and Construction Summary

Capital Operating Expenses

(dollars in thousands)

	FY 2004	FY 2005	FY 2006	$ Change	% Change
Capital Equipment	9,185	8,000	6,000	-2,000	-25.0%