Advanced Scientific Computing Research

Funding Profile by Subprogram

(dollars in thousands)

	FY 2008 Current Appropriation	FY 2009 Original Appropriation	FY 2009 Additional Appropriation[a]	FY 2010 Request
Advanced Scientific Computing Research				
Mathematical, Computational, and Computer Sciences Research	117,900	144,205	+37,130	163,792
High Performance Computing and Network Facilities	223,874	224,615	+119,980	245,208
Total, Advanced Scientific Computing Research	341,774[b]	368,820	+157,110	409,000

Public Law Authorizations:
Public Law 95–91, "Department of Energy Organization Act", 1977
Public Law 108–423, "Department of Energy High-End Computing Revitalization Act of 2004"
Public Law 109–58, "Energy Policy Act of 2005"
Public Law 110–69, "America COMPETES Act of 2007"

Program Overview

Mission

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy. A particular challenge of this program is fulfilling the science potential of emerging multi-core computing systems and other novel "extreme-scale" computing architectures, which will require significant modifications to today's tools and techniques.

Background

Imagine exploring the inner workings of a supernova or traveling backward and forward in time to observe changes in our global climate. Scientists today are able to explore these realms thanks to a 100 fold increase in computing power delivered over the past five years and to the software developed to harness the power of these forefront computers.

Throughout recorded history, humankind has always been interested in understanding the mysteries of the universe. Mathematics has been the primary tool. Pythagoras used mathematics to determine the diameter of the earth. Newton and Leibniz invented calculus to understand the movement of the planets. Mathematical research in the 1800s led to Einstein's Theory of General Relativity.

While mathematics enabled the study of increasingly complex problems, the time to carry out these calculations became unmanageable. Today, advances in mathematics and computing enable scientists to understand everything from Alzheimer's disease to climate change. ASCR and its predecessor organizations have led these advances for the past thirty years.

[a] The Additional Appropriation column reflects the planned allocation of funding from the American Recovery and Reinvestment Act of 2009, P.L. 111–5. See the Department of Energy Recover website at http://www.energy.gov/recovery for up-to-date information regarding Recovery Act Funding.
[b] Total is reduced by $9,399,000: $8,392,000 of which was transferred to the Small Business Innovative Research (SBIR) program and $1,007,000 of which was transferred to the Small Business Technology Transfer (STTR) program.

In past year, ASCR has delivered petascale computing power to the open science community and two of the world's fastest computers for open-science. These machines enabled two Gordon Bell prizes, including the world's first petascale application, and swept the High Performance Computing challenge at SuperComputing08. Since the machines are open to all on a competitive basis, the scientific applications that are running provide new insights into areas such as: first-principles flame simulation to guide design of fuel-efficient clean engines, high temperature superconductivity, the molecular basis for Parkinson's disease, supernova shock wave instability, designing proteins at atomic scale and creating enzymes.

This growing area of research allows scientists to see things that can't be seen experimentally and to reveal new insights from experimental data. Computational scientists create mathematical models and simulations of physical, biological and chemical phenomena, which allow them to better understand the phenomena and to predict behavior. In the case of climate change, we cannot afford to simply wait and record the impacts of increasing concentrations of greenhouse gases. Through computation, scientists can model what is known about the Earth system and identify experiments to gather necessary data to further refine the model. The model allows us to predict, with increasing confidence, future changes in the climate and identity mitigation strategies for policy makers. Such research has made computational science a true third pillar of science, alongside theory and experiment.

ASCR supports basic research in both applied mathematics and computer science focused in areas relevant to high performance computing. It links this research to scientists across SC through the Scientific Discovery through Advanced Computing (SciDAC) program. SciDAC facilitates transfer of the results of basic efforts into computational science through direct partnerships between applied mathematicians, computer scientists and domain experts in a specific discipline. These partnerships have been spectacularly successful with documented improvements in code performance in excess of 10,000 percent.

The other primary goal of the ASCR program is to remove geography as a barrier to science. Even before the Internet, researchers across the Office of Science have been communicating with each other, exchanging large data sets and running complex calculations and experiments in remote scientific user facilities. ASCR has had a leading role in driving development of the networks connecting these researchers. Even today, the invisible glue that binds all the networks in the world together and effortlessly passes billions of searches and trillions of bits has roots in ASCR research. ASCR researchers helped to establish critical protocols such as TCP/IP on which the current Internet is based. ASCR advanced networking research also makes international collaborations, such as the Large Hadron Collider and ITER, possible. The Internet has removed barriers between people and ASCR's advanced scientific networks have removed barriers between scientists and research facilities.

Looking forward, ASCR must continue to take its computers and networks to the leading edge of technology and transform them into tools for scientific discovery. Therefore, ASCR must address the emerging challenges of next generation computing systems and transforming extreme scale data into knowledge. The ASCR approach of integrating research results across disciples and with forefront facilities has been the key to our history of success in computational science. With this integrated approach ASCR will continue to deliver scientific insight to address national problems.

Strategic and GPRA Unit Program Goals

The ASCR program has one Government Performance and Results Act (GPRA) Unit Program Goal which contributes to Strategic Goals 3.1 and 3.2 in the "goal cascade":

- GPRA Unit Program Goal 3.1/2.51.00: 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 GPRA Unit Program Goal 3.1/2.51.00, Deliver forefront computational and networking capabilities

ASCR contributes to this goal 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 both facilities and colleagues to enable scientific discovery.

ASCR supports fundamental research in applied mathematics, computer science, networking research, 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 SciDAC; and provides the advanced computing and network resources that enable scientists to use these tools for scientific inquiry. Applied mathematics enables scientists to accurately model physical and natural systems, 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 advance the frontiers of simulation and scientific discovery. Shrinking the distance between scientists and the resources they need is also critical to 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 networking research provide the capability to move the millions of gigabytes that these resources generate to the scientists' desktops. Therefore, the ASCR program contributes to research programs across SC, as well as other elements of the Department.

The following indicators establish specific long term (ten year) goals in Scientific Advancement that the ASCR program is committed to, and progress can be measured against. The Advanced Scientific Computing Advisory Committee (ASCAC) was charged to review progress toward these long term measures and reported "good to excellent" progress to the Department in November, 2006. The long term measures are:

- Develop multiscale mathematics, numerical algorithms, and software that enable more 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; and
- 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 2005 Results	FY 2006 Results	FY 2007 Results	FY 2008 Results	FY 2009 Targets	FY 2010 Targets
GPRA Unit Program Goal 3.1/2.51.00 (Deliver Forefront Computational and Networking Capabilities)					
Mathematical, Computational, and Computer Sciences Research					
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]	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. FY 2006—>50%. [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. FY 2007—>100% [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. FY 2008—>100% [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. FY 2009—>100%	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, tools or libraries. FY 2010—>100%
High Performance Computing and Network Facilities					
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]					
Focused usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used that was accounted for by computations that require at least 1/8 of the total resource. [Goal Met]	Focused usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used that was accounted for by computations that require at least 1/8 of the total resource. FY 2006—40%. [Goal Met]	Focused usage of the primary supercomputer at the NERSC on capability computing. Percentage of the computing time used that was accounted for by computations that require at least 1/8 (760 processors) of the total resource. FY 2007—40% [Goal Met]	Focus usage of the primary supercomputer at the National Energy Research Scientific Computing Center (NERSC) on capability computing. Thirty percent (30%) of the computing time will be used by computations that require at least 1/8 (2,040 processors) of the NERSC resource. FY 2008 goal 30%. [Goal Met]	Focus usage of the primary supercomputer at the National Energy Research Scientific Computing Center (NERSC) on capability computing. Thirty percent (30%) of the computing time will be used by computations that require at least 1/8 (2,040 processors) of the NERSC resource. FY 2009 goal 30%.	Focus usage of the primary supercomputer at the National Energy Research Scientific Computing Center (NERSC) on capability computing. Forty percent (40%) of the computing time will be used by computations that require at least 1/8 (2,040 processors) of the NERSC resource. FY 2010 goal 40%.

Subprograms

To accomplish its mission and address the challenges described above, the ASCR program is organized into two subprograms—Mathematical, Computational, and Computer Sciences Research and High Performance Computing and Network Facilities.

- The *Mathematical, Computational, and Computer Sciences Research* subprogram develops mathematical descriptions, models, methods, and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales. Examples include the behavior of the earth's climate, living cells, and the reactive transport of contaminants through groundwater. The subprogram also develops the software to make effective use of advanced networks and computers, many of which contain thousands of multi-core processors with complicated interconnections, and to transform enormous data sets from experiments and simulations into scientific insight. In addition, the subprogram supports computational science partnerships within the Office of Science, R&D integration efforts with the Department's applied programs, and interagency collaborations.
- The *High Performance Computing and Network Facilities* subprogram delivers forefront computational and networking capabilities and contributes to the development of next-generation capabilities through support of prototypes and testbeds.

Effective scientific utilization of high-end capability computing requires dynamic partnerships among application scientists, applied mathematics, computer scientists and facility support staff. Therefore, close coordination both within and across ASCR subprograms and with partner organizations is key to the success of the ASCR program.

Benefits

Computer-based simulation enables us to model the behavior of complex systems that are beyond the reach of experiment or for which there is no theory. Because computer-based simulation is so important to research programs across the Office of Science and throughout the government, the ASCR Leadership Computing Facilities are operated as open user facilities, with access determined by merit evaluation of proposals.

Many of the applications running on these facilities have direct benefit to science and society at large.

- Computational chemistry and simulation of nanomaterials is relevant to energy applications. These applications are funded in partnership with Basic Energy Sciences.
- The next generation Earth System Models will dramatically improve our ability to predict changes in global climate. This work is funded in partnership with Biological and Environmental Research. ASCR also provides the majority of the computing and networking resources for the U.S. contributions to the Intergovernmental Panel on Climate Change.
- Simulations of fusion reactors help to develop fusion as a clean, abundant energy source. This work is jointly funded with the Office of Fusion Energy Sciences.
- Computer modeling of nuclear structure has relevance for science, nuclear energy, and nuclear weapons. These applications are through partnerships with both Nuclear Physics and the National Nuclear Security Administration (NNSA).
- Understanding the mysteries of the universe is the focus of partnerships with High Energy Physics that include experiments, such as three dimensional simulations of supernovae events, which are only possible with leadership computing resources.

- Computational biology is relevant for energy and bioremediation applications. This work is a partnership with the Biological and Environmental Research Office. Projects from other agencies have investigated proteins; blood flow, and Parkinson's disease.
- Subsurface science research characterizes and predicts changes in the Department's environmental management sites. This work also has implications for the Department's efforts in subsurface carbon sequestration. These applications are partnerships with the Biological and Environmental Research.

Establishing leadership computing for the Office of Science required partnering with the hardware vendors to develop the software necessary to effectively use these powerful systems. These partnerships benefit many sectors of the economy from high-tech industry and academic research to software development and engineering.

Finally, ASCR's support of researchers and students (the next generation of researchers) is a benefit to the national research and development workforce.

Program Planning and Management

The ASCR program has developed a system of planning and priority setting that relies heavily on input from groups of outside experts. ASCR has also instituted a number of peer review and oversight processes designed to assess the quality, relevance, and performance of the ASCR portfolio on a regular basis. ASCR peer reviews all of its activities. One way in which ASCR ensures the integrity and effectiveness of the peer review processes is through the Advanced Scientific Computing Advisory Committee (ASCAC), which organizes regular Committees of Visitors to review ASCR research management, reviews output of the ASCR scientific user facilities, and reviews progress toward the long-term goals of the program. In addition, ASCAC identifies scientific challenges and opportunities, including specific bottlenecks to progress in areas such as climate change or computational biology, and comments on the overall balance of the ASCR portfolio.

In addition to ASCAC, critical tools for managing the ASCR scientific user facilities are annual operational reviews and requirements workshops. For example, ESnet conducts two network requirements workshops per year with individual DOE Office of Science program offices. The purpose of each workshop is to accurately characterize the near-term, medium-term and long-term network requirements of the science conducted by each program office. Since two workshops are conducted per year, ASCR refreshes the network requirements information for each of the six program offices every three years.

Community driven workshops are another critical means by which dialogues are facilitated and new research opportunities are identified. For example, ASCR is sponsoring a series of workshops in 2009, in partnership with the other Office of Science programs and the NNSA, to identify key science challenges in the disciplines important to the Department and the potential role of extreme scale computing in addressing those challenges.

Another important planning and coordination mechanism for ASCR is the National Science and Technology Council's (NSTC) subcommittee on Networking and Information Technology Research and Development (NITRD). ASCR is a major participant in the NITRD Program[a] which coordinates Federal research investments by the 11 member agencies in advanced information technologies such as computing, networking, and software through interagency working groups and coordinating groups. ASCR is a major participant and/or chair of the High End Computing Research and Development, Large Scale Networking, and High End Computing Infrastructure and Applications groups. In FY 2008, a

[a] Information on the NITRD Program can be found at http://www.nitrd.gov.

NITRD interagency working group, led by ASCR, the Department of Defense and the National Science Foundation, developed the "Federal Plan for Advanced Networking Research and Development"[a] to identify key research areas necessary to deliver future networking technologies that are critical for science but also for U.S. economic and national security. This interagency plan provides the framework in which ASCR research will address key issues for science and effectively leverage the research efforts of other agencies.

In October 2008, the National Research Council published a study titled "The Potential Impact of High-End Capability Computing (HECC) on Four Illustrative Fields of Science and Engineering"[b] that identifies and categorizes important scientific questions and technology problems for which an extraordinary advancement in our understanding is difficult or impossible without leading edge scientific simulation capabilities. In all four fields studied—atmospheric sciences, astrophysics, separations chemistry, and evolutionary biology—the committee found continuing demand for more powerful HECC and for large scale data management. The report outlined the major scientific challenges in the four fields and estimated the associated challenges in mathematics, computer science, and computing infrastructure. The conclusions of the report underscore the importance of balancing investments in high potential application areas, the high-end computing resources required by multiple fields, and the longer-term mathematics and computer science research that underpins continued progress. However, the report also emphasizes the added importance of linking these efforts: "In many cases HECC capabilities must continue to be advanced to maximize the value of data already collected…The committee foresees a growing need for computational scientists and engineers who can work with mathematicians and computer scientists to develop next-generation code."[c]

Basic and Applied R&D Coordination

A cornerstone of the ASCR program is coordination across disciples and programs. Partnerships within the Office of Science are mature and continue to advance use of high performance computing and scientific networks for science. In addition, ASCR continues to have a strong partnership with NNSA in areas of mutual interest including best practices for management of high performance computing facilities. Through NITRD, ASCR coordinates with similar programs across the federal government and directly partners with the Department of Defense on developing High Productivity Computing Systems and software and with the National Science Foundation on the Open Science Grid.

In discussions with the applied R&D programs in the Department, a key area of mutual interest continues to be in applied mathematics for optimization of complex systems, control theory, and risk assessment. This was the subject of an ASCR workshop held in December 2006 that identified research challenges in advanced mathematics that could benefit the optimization of fossil fuel power generation, the nuclear fuel lifecycle, and power grid control. Such research could increase the likelihood for success in DOE strategic initiatives such as smart grid. Another workshop, specifically focused on the challenges of grid modernization efforts, was organized in partnership with the Office of Electricity Delivery and Energy Reliability in March 2009. This workshop is part of a series of workshops on basic research needs in applied R&D areas. Other workshops have covered advanced nuclear energy systems (with the Office of Nuclear Energy), subsurface science (with the Office of Environmental Management, the Office of Fossil Energy, and the Office of Civilian Radioactive Waste Management), cyber security

---

[a] The "Federal Plan for Advanced Networking Research and Development" can be found at http://www.nitrd.gov/Pubs/ITFAN-FINAL.pdf

[b] The "The Potential Impact of High-End Capability Computing (HECC) on Four Illustrative Fields of Science and Engineering" can be founds at http://www.nap.edu/catalog.php?record_id=12451

[c] "The Potential Impact of High-End Capability Computing (HECC) on Four Illustrative Fields of Science and Engineering", page 7–8

(with the Office of Electricity Delivery and Energy Reliability), and alternative and renewable energy (with the Office of Energy Efficiency and Renewable Energy). These workshops facilitate a dialogue between the ASCR research community and a specific applied R&D research community and identify opportunities for new research. This research becomes part of the ASCR program through investigator driven research proposals and is coordinated with the applied efforts through program manager interactions and joint principal investigator meetings.

Budget Overview

The FY 2010 ASCR budget request capitalizes on the significant gains in computational science over the past decade and positions the Department to attack scientific challenges through modeling and simulation in the next decade. The request also balances investments in high performance computing facilities and advanced networks with investments in applied mathematics, computer science, next generation networks for science and computational partnerships. This balance ensures continued progress in a wide array of fields important to the Department's missions for FY 2010 and for years to come.

The FY 2010 budget request continues support for the Leadership Computing Facility (LCF) at Oak Ridge National Laboratory (OLCF)—the most capable machine in the U.S. for open science, a 1.64 petaflop, multicore Cray Baker system—and will make that system openly available to the scientific community through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. In addition to INCITE projects, the OLCF will continue to provide access and assistance to tool and library developers and to researchers seeking to scale their application to this new realm of computing power. These activities are critical to harnessing the complexity of this architecture and to respond to the challenges associated with high performance computer systems with large numbers of multicore processors. This effort is also expected to build experience and tools for the synergistic DARPA High Productivity Computing Systems award to Cray, supported in part through the Research and Evaluation Prototypes activity, and other anticipated architectures that will exhibit even greater complexity.

The Argonne LCF (ALCF) will continue allocating the 556 teraflop IBM Blue Gene/P through INCITE in FY 2010. The FY 2010 budget request also includes site preparation activities for a next generation machine such as the one being developed through the joint research project with NNSA, IBM, Argonne and Lawrence Livermore National Laboratories—an IBM Blue Gene/Q—that is supported through the Research and Evaluations Prototypes activity.

The National Energy Research Scientific Computing (NERSC) facility at Lawrence Berkeley National Laboratory (LBNL) will operate at a capacity of over 356 teraflops in FY 2010 and will provide the production computing resources for the Office of Science programs. Focus will be on assisting applications to effectively utilize the potential of this facility and to move beyond NERSC to the larger LCF machines. Acquisition of NERSC 6 in FY 2010 will bring total NERSC capacity to approximately one petaflop to meet ever growing demand from Office of Science researchers.

The FY 2010 budget request supports ESnet to continue to advance the next generation network capability that is critical to Department applications and facilities. ESnet will deliver 100–400 gigabits per second (Gbs) connections to Office of Science laboratories in FY 2010, with a goal of achieving 1,000 Gbs connectivity in 2014. These increases in bandwidth are necessary to move massive amounts of data to and from the petascale computing facilities and from other research facilities such as the Large Hadron Collider and Spallation Neutron Source. The ESnet is also critical to effective utilization of the growing amounts of data in climate research, nuclear structure, genomics, and proteomics that advance the Department's energy and environment missions.

The FY 2010 budget continues the research efforts in Scientific Discovery through Advanced Computing (SciDAC) and the core research programs in Applied Mathematics and Computer Science that enable scientists to effectively utilize the capabilities of the LCFs while beginning to lay the basic research foundation necessary to realize the potential from the more complex systems on the horizon. Increases in core research in Applied Mathematics and Computer Science for FY 2010 will be targeted at long-term research needs. In networking, the growing amounts of scientific data and the need to remotely access large-scale scientific facilities will challenge the current network infrastructure. New technologies are emerging to enhance the capacity of the existing network but new techniques and tools will need to be developed to effectively utilize that capacity. The Next Generation Networking for Science research portfolio was re-competed in FY 2009 to address the most critical research needs for these next generation scientific networks.

Significant Program Shifts

A new allocation option for the ASCR scientific computing facilities was established in FY 2009 called the ASCR Leadership Computing Challenge program. This program is open year-round to scientists from the research community in academia and industry and allocates up to 30 percent of the computational resources at NERSC and the LCFs at Argonne and Oak Ridge for special situations of interest to the Department with an emphasis on high-risk, high-payoff simulations in areas directly related to the Department's energy mission, for national emergencies, or for broadening the community of researchers capable of using leadership computing resources. The majority of the LCF resources will continue to be allocated through INCITE. NERSC will be focused on the program priorities of the Office of Science.

Mathematical, Computational, and Computer Sciences Research

Funding Schedule by Activity

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
Mathematical, Computational, and Computer Sciences Research			
Applied Mathematics	32,146	40,164	44,850
Computer Science	24,121	33,618	46,800
Computational Partnerships	54,053	52,064	53,235
Next Generation Networking for Science	7,580	14,321	14,321
SBIR/STTR	—	4,038	4,586
Total, Mathematical, Computational, and Computer Sciences Research	117,900	144,205	163,792

Description

The Mathematical, Computational, and Computer Sciences Research (Research) subprogram supports the research elements of the ASCR program to effectively utilize forefront computational and networking capabilities. Computational science is increasingly central to progress at the frontiers of science and to our most challenging feats of engineering. Accordingly, the Research subprogram must be positioned to address scientifically challenging questions, such as:

- What new mathematics are required to accurately 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?
- What operating systems, data management, analysis, representation model development, user interface, and other tools are required to make effective use of future-generation supercomputers?
- What tools are needed to make all scientific resources readily available to scientists, regardless of whether they are at a university, national laboratory, or industrial setting?

FY 2008 Accomplishments

- *SciDAC team wins 2008 Gordon Bell special prize for algorithms*. The SciDAC team won for developing the Linearly Scaling 3D Fragment (LS3DF) Method, which was used to predict the energy harnessing efficiency of nanostructures that can be used in solar cell design. The team computed the electronic structure of a 3,500-atom ZnTeO (zinc-telluride-oxygen) alloy to verify that the code could be used to predict properties of the alloy that had previously been experimentally observed. The simulation led to a prediction for the efficiency of this alloy as a new solar cell material. The prestigious Gordon Bell special prize was awarded for the clever way in which the code reduces the complexity of the problem to something within reach of today's computers without sacrificing scientific content.

- *First-Principles flame simulation provides crucial information to guide design of fuel-efficient cleaner engines.* Combustion models for designing tomorrow's vehicle engines and power generation devices challenge the capabilities of the most advanced computers but provide valuable insights toward improving efficiency and reducing emissions. These researchers managed to model flames with a wide-range of characteristics. They found that if low-temperature compression ignition concepts employing lean dilute fuel mixes are widely adopted in next generation engines, fuel efficiency could increase by as much as 25-50%. This approach would also help meet future low-emission vehicle standards with almost undetectable emissions of nitrogen oxide, a major contributor to smog.
- *First provably scalable solver for Maxwell's equations*. A team of applied mathematicians has developed the first provably scalable solver code for Maxwell's equations, a set of equations that are fundamental to physics and engineering of electric and magnetic fields. This novel software technology, known as an auxiliary-space Maxwell solver, outperforms earlier solution techniques by as much as 25 times. This enables researchers to solve large-scale computational electromagnetics problems with far greater accuracy. Electromagnetic simulations have a wide range of applications such as in the development of pulsed-power devices, semiconductor chips, stealth aircraft, and electrical generators.
- *Powerful mathematical tools resolve complex simulations.* PETSc (Portable, Extensible Toolkit for Scientific computation) was developed as a general purpose suite of tools for solving equations used in large-scale application projects. The PETSc solvers are used to model complex phenomena in virtually all areas of DOE-sponsored science and engineering research, including climate science, fission and fusion energy, nanoscience simulations, subsurface flows, oil-reservoir modeling, combustion, fracture mechanics, and micromagnetics. A 2009 example of an application using PETSc is PFLOTRAN, a next generation reactive flow and transport model that provided important information about reducing uranium concentrations at the Hanford 300 site.
- *Astrophysicists discover supernova shock-wave instability and a better way to spin up pulsars.* SciDAC researchers have provided two key pieces in the puzzle of the core-collapse supernova, demonstrating that the supernova shock wave is unstable and showing how the shock wave instability may be responsible for the spin of the leftover pulsar. Their efforts help us understand how some of the universe's most dramatic catastrophes are responsible for producing and spreading the Elements of life.
- *SciDAC accelerator project influences accelerator design and addresses construction challenges at DOE facilities.* The Community Petascale Project for Accelerator Science and Simulation (ComPASS) project, in collaboration with one of the math-focused Centers for Enabling Technologies, developed an uncertainty quantification tool to reconstruct the cavity shape of a particle accelerator using measured data as inputs. The tool successfully identified the cause of the beam breakup instability in a cryomodule of the 12 GeV CEBAF Upgrade Project at the Thomas Jefferson National Accelerator Facility and the findings helped the project move into construction.

Detailed Justification

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
Applied Mathematics	**32,146**	**40,164**	**44,850**

The Applied Mathematics activity supports the research, development, and application of applied mathematical models, methods and algorithms to understand complex physical, chemical, biological, and engineered systems related to the Department's mission. The research falls into eight general categories described below. The first five have been supported for a number of years and the last three began in FY 2009.

- Numerical methods research for equations related to problems such as wave propagation, electrodynamics, fluid flow, elasticity, and other natural or physical processes.
- Advanced linear algebra research for fast and efficient numerical solutions of linear algebraic equations that often arise when simulating physical processes. Because a large fraction of the time in many simulations is spent doing this type of computation, advances here have enormous leverage across science.
- Computational meshing research for developing ways in which space can be broken up into regions—often geometrically complex—for the purposes of simulation.
- Optimization research for mathematical methods for minimizing energy or cost, or finding the most efficient solutions to engineering problems, or discovering physical properties and biological configurations. This includes optimization, control, and risk assessment in complex systems with relevance for DOE missions in energy, national security, and environment.
- Multiscale mathematics and multiphysics computations for connecting the very large with the very small, the very long with the very short, and multiple physical models in a single simulation.
- Joint Applied Mathematics-Computer Science Institutes for the development of efficient new mathematical models, algorithms, libraries, and tools for next generation computers.
- Mathematics for the analysis of extremely large datasets for identifying key features, determining relationships between the key features, and extracting scientific insights.
- Mathematics of cyber security from a basic research perspective for addressing the understanding and discovery of anomalies in existing network data, modeling of large-scale networks, and understanding dynamics and emergent behavior on networks. This leverages on-going efforts in the mathematics of optimization and risk assessment in complex systems.

These mathematical models, methods and algorithms are the fundamental building blocks for describing physical systems computationally. Applied Mathematics research underpins all of the modeling and simulation efforts in the Department.

In FY 2010, the Applied Mathematics activity will support the long term cyber security challenges of open science that was transferred from Next Generation Networking for Science and initiated in FY 2009. Early career awards given to exceptionally talented university investigators in Applied Mathematics will also continue to be supported. The Computational Science Graduate Fellowship program, aimed at attracting the best graduate students in the scientific disciplines and educating them

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010

as the next generation of computational scientists, is continued at $6,000,000. In addition, in FY 2010 the Applied Mathematics activity will initiate a new fellowship program for graduate students and young investigators in applied mathematics and high performance computer science.

Computer Science	**24,121**	**33,618**	**46,800**

The Computer Science activity supports research to utilize computing at extreme scales and to understand extreme scale data from both simulations and experiments. The research falls into five general categories described below. The first four have been supported for a number of years and the last was begun in FY 2009.

- Operating and file systems for extreme scale computers with many thousands of multi-core processors with complicated interconnection networks.
- Performance and productivity tools for extreme scale systems that enable users to diagnose and monitor the performance of software and scientific application codes to enable users to improve performance and get scientific results faster.
- Programming models that enable today's computations and discover new models that scale to hundreds of thousands of processors to simplify application code development for petascale computing.
- Data management and visualization: to transform extreme scale data into scientific insight through investments in visualization tools that scale to multi-petabyte datasets and innovative approaches to indexing and querying data.
- Joint Applied Mathematics-Computer Science Institutes for the development of efficient new mathematical models, algorithms, libraries, and tools for next generation computers. Leading edge developers to directly address the new challenges from the next generation of computers and transfer this insight to key DOE application developers.

The Computer Science activity addresses two fundamental challenges. The first challenge is enabling science applications to harness computer systems with increasing scale and increasing complexity due to technology advances such as multicore chips. This challenge will require more dynamic behavior of system software (operating systems, file systems, compilers, performance tools) than historically developed. Substantial innovation is needed to provide essential system software functionality in a timeframe consistent with the anticipated availability of hardware. The second challenge is enabling scientists to effectively manage, analyze and visualize the petabytes of data that result from extreme scale simulations and experimental facilities. Substantial innovation in computer science and applied mathematics is needed to provide essential system and application functionality in a timeframe consistent with the anticipated availability of hardware.

In FY 2010, the Computer Science activity will initiate a new effort in advanced computer architecture design for science to ensure that ASCR will be actively engaged in the early stages of continued advancement in high performance computing systems suitable for DOE's scientific applications. This new effort will focus on pre-competitive research that falls short of the fabrication and testing of prototypes, and will involve close coordination with other agencies and high performance computing

	(dollars in thousands)	
FY 2008	FY 2009	FY 2010

vendors. In addition, the Computer Science activity will expand support for the most promising long-term research to address the challenges of next generation computing systems and extreme scale data.

This research will bridge efforts in advanced computer architecture design with ongoing efforts in computer science and applied mathematics to address the specific needs of DOE science applications. Early career to exceptionally talented university investigators in Computer Science will also continue.

	FY 2008	FY 2009	FY 2010
Computational Partnerships	**54,053**	**52,064**	**53,235**

The Computational Partnership activity supports the Scientific Discovery through Advanced Computing (SciDAC) program to use results from applied mathematics and computer science research on scientific applications sponsored by other Office of Science programs. These partnerships enable improved performance on high-end computing systems for scientists to conduct complex scientific and engineering computations at a level of detail that begin to resemble real-world conditions. The activities fall into four general categories described below.

- The nine multi-institutional *Centers for Enabling Technologies* represent almost half of the ASCR SciDAC activity. They are a focal point for bringing together a critical mass of leading experts from multiple disciplines to focus on key problems in a particular area such as performance, data management, optimization, or visualization. These SciDAC Centers address needs for new methods, algorithms and libraries; new methodologies for achieving portability and interoperability of complex scientific software packages; software tools and support for application performance; and more effective tools for feature identification, data management, and visualization.
- The four multi-institutional *SciDAC Institutes* are university-led centers of excellence which complement the efforts of the SciDAC Centers for Enabling Technologies but with a role in the education and training of the next generation of computational scientists. These institutes reach out to a broader community of scientists to advance scientific discovery through advanced computation, collaboration, and training of graduate students and postdoctoral fellows.
- The 35 multi-institutional *Science Applications Partnerships* are partnerships with other Office of Science programs to dramatically improve the ability of their researchers to effectively utilize petascale computing to advance science. These partnerships support collaborative research between applied mathematicians and computer scientists (supported by ASCR) with domain scientists (supported by the other programs) to refine and apply computational techniques and tools that address the specific problems of a particular research effort, such as modeling the reactive transport of contaminants through groundwater or developing an Earth System model that fully simulates the coupling between the physical, chemical, and biogeochemical processes in the global climate.
- With more than 80 participating institutions and hundreds of researchers developing tools, techniques, and software that push the state-of-the-art in high performance computing, ASCR needed to ensure that SciDAC teams shared information across projects and, to leverage taxpayer investment, with other researchers. The *SciDAC Outreach Center*, a small virtual organization linked to user support at the ASCR facilities and organized by LBNL, provides a single resource to facilitate and accelerate the transfer of tools, techniques, and methods to the broader research community. The SciDAC Outreach Center is also a resource for INCITE applicants who need assistance in readying their application for leadership facilities.

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010

In FY 2010, the Computational Partnerships activity will support a small number of new interdisciplinary teams focused on transforming critical DOE applications for extreme scale computing. These competitively selected teams will evaluate the impact of directions in computer hardware on application capability, form a critical interface to existing SciDAC Centers and Institutes on the tool and library implications of these developments and develop the understanding needed to enable these applications to execute effectively on future computer architectures.

Next Generation Networking Research for Science	**7,580**	**14,321**	**14,321**

The Next Generation Networking for Science activity builds on results from Computer Science and Applied Mathematics to develop integrated software tools and advanced network services to enable large-scale scientific collaboration and to utilize the new capabilities of ESnet to advance DOE missions. The research falls into two general categories described below.

- Distributed systems software including scalable and secure tools and services to facilitate large-scale national and international scientific collaboration and high-performance software stacks to enable the discovery, management, and distribution of extremely large data sets generated by simulations or by science experiments such the Large Hadron Collider, the Intergovernmental Panel on Climate Change, and ITER.
- Advanced network technologies including dynamic optical network services, scalable cyber security technologies, and multi-domain, multi-architecture performance protocols to seamlessly interconnect and provide access to distributed computing resources and science facilities.

In FY 2010, Next Generation Networking for Science activity transfers long term cyber security research to the Applied Mathematics activity and will initiate new research efforts to focus on developing technologies to support research and education networks such as ESnet. Early career awards given to exceptionally talented university investigators in Networking will also continue to be supported.

Small Business Innovative Research (SBIR)/ Small Business Technology Transfer (STTR)	**—**	**4,038**	**4,586**

In FY 2008, \$3,354,000 and \$403,000 were transferred to the SBIR and STTR programs respectively. The FY 2009 and FY 2010 amounts shown are the estimated requirements for continuation of the congressionally mandated SBIR and STTR programs.

Total, Mathematical, Computational, and Computer Sciences Research	**117,900**	**144,205**	**163,792**

Explanation of Funding Changes

	FY 2010 vs. FY 2009 ($000)
Applied Mathematics	
This increase reflects the shift of cyber security research from Next Generation Networking for Science to Applied Mathematics and supports a new fellowship program for graduate students and young investigators in applied mathematics and high performance computer science.	+4,686
Computer Science	
This increase will support a new effort in advanced computer architecture design for science and will expand support for the most promising long-term research to address the challenges of next generation computing systems and extreme scale data. This research will bridge efforts in advanced computer architecture design with ongoing efforts in computer science and applied mathematics to address the specific needs of DOE science applications.	+13,182
Computational Partnerships	
The increase will support a small number of interdisciplinary teams focused on transforming critical DOE applications for extreme scale computing.	+1,171
SBIR/STTR	
Increase in SBIR/STTR due to increase in operating expenses.	+548
Total Funding Change, Mathematical, Computational, and Computer Sciences Research	**+19,587**

High Performance Computing and Network Facilities

Funding Schedule by Activity

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
High Performance Computing and Network Facilities			
High Performance Production Computing	56,700	54,790	55,000
Leadership Computing Facilities	114,065	115,000	130,000
Research and Evaluation Prototypes	25,347	23,900	23,900
High Performance Network Facilities and Testbeds	27,762	25,000	29,862
SBIR/STTR	—	5,925	6,446
Total, High Performance Computing and Network Facilities	223,874	224,615	245,208

Description

The High Performance Computing and Network Facilities (Facilities) subprogram delivers forefront computational and networking capabilities to scientists nationwide.

To maintain leadership in areas of scientific modeling and simulation important to DOE missions, the Facilities subprogram plans, develops, and operates high performance computing facilities and advanced networks that are available 24 hours a day, 365 days a year. This includes the High Performance Production Computing at the National Energy Research Scientific Computing (NERSC) facility at LBNL, Leadership Computing Facilities (LCFs) at Oak Ridge and Argonne National Laboratories, and High Performance Network Facilities and Testbeds through the Energy Sciences Network (ESnet) managed by LBNL. The Facilities subprogram also invests in long-term needs through the Research and Evaluation Prototypes activity.

The Facilities subprogram contributes to DOE missions by providing the leadership and high performance computing facilities and scientific networks that support mission driven and open science research. These computers, and the other SC research facilities, turn out many petabytes of data each year. Moving these data to the researchers who need them requires advanced scientific networks and related technologies also provided through the Facilities subprogram.

The Facilities subprogram computing resources are allocated through competitive processes. The LCFs are predominately allocated through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program to researchers at universities, national laboratories, and foreign institutions. Up to eighty percent of the processor time on LCFs is allocated through INCITE to a small number of projects, each requiring a substantial amount of the available resources. LCFs provided over 889 million hours of computing time in calendar year 2009. The high performance production computing facilities at NERSC are focused on the computing needs of the SC and allocated through a competitive process reserved for researchers supported by the SC programs. In FY 2010, all of the ASCR scientific computing facilities will allocate up to thirty percent of computing resources through the ASCR Leadership Computing Challenge program as described under Significant Program Shifts.

FY 2008 Accomplishments

- *World's first petascale application provides new insights into superconductivity.* Researchers simulated chemical disorder in high-temperature superconductors known as cuprates—layers of copper oxide separated by layers of an insulating material. These simulations were the first where there was enough computing power to move beyond ideal, perfectly ordered materials. By advancing understanding of the interplay between these imperfections and superconductivity, the work promises to help researchers push transition temperatures ever higher, possibly approaching the goal of room-temperature superconductivity, or materials that exhibit this behavior without artificial cooling. This work received the prestigious 2008 Association for Computing Machinery Gordon Bell Prize for attaining the fastest performance ever in a scientific supercomputing application—1.352 quadrillion calculations a second, or 1.352 petaflops—on the Oak Ridge LCF.
- *Oak Ridge leadership computing facility world's first petaflop computer dedicated to open science.* The latest upgrade to the LCF increases the system's computing power to a peak 1.64 petaflops, or quadrillion mathematical calculations per second. The project to build a petaflops machine—completed on time, on budget and exceeding the original scope—included partnerships with industry to develop new hardware and computer architectures. The LCF is made available to the scientific community through SC's peer-reviewed INCITE program. In 2008, INCITE research included accelerator physics, astrophysics, chemical sciences, climate research, computer science, engineering physics, environmental science, fusion energy, life sciences, materials science, nuclear physics, and nuclear engineering. Practical applications of the research include improving commercial aircraft design, advancing fusion energy, studying supernova, understanding nanomaterials, and studying global climate change and the causes of Parkinson's disease.
- *National Energy Research Scientific Computing (NERSC) facility research advances goals for a green economy.* Researchers investigating the renewable production and storage of hydrogen at NERSC have uncovered ways of reducing carbon dioxide using methods that mimic photosynthesis. These results open new opportunities for photochemical reactions by visible-light irradiation and was the cover article for the May 2008 issue of Inorganic Chemistry.
- *Argonne Leadership Computing Facility (ALCF) uses an innovative cooling system to reduce power consumption.* Among the top 20 supercomputers in the world in November, 2008, the ALCF is the second-most energy-efficient and needs only a little more than one megawatt of power. Much of the electricity that the ALCF requires is not actually used to process computations but rather to cool machinery. Without any cooling at all, the room that houses the computer and peripherals would reach 100 degrees Fahrenheit in ten minutes after the equipment is powered on. To reduce the amount of electricity used, the ALCF was able to take advantage of Chicago's cold winters and replaced air conditioners with fans to move 300,000 cubic feet of water-chilled air per minute to maintain a room temperature of 64 degrees Fahrenheit. The technique uses only 60 percent more energy for cooling than the supercomputer itself draws, compared to over twice the power for typical supercomputers. In addition, the ALCF's chilled water plant uses cooling towers to chill the water when the weather is cold. Once the temperature falls to 35 degrees Fahrenheit or below outside, the temperature in the chilled water system is maintained solely by Argonne's cooling towers saving about $20,000 to $25,000 a month in electrical costs.
- *Energy Sciences Network (ESnet) launches first segment of the next-generation nationwide scientific research network.* The Department of Energy completed upgrades to ESnet4 in December 2008. ESnet4 is one of the most advanced and reliable, high capacity nationwide networks supporting

scientific research. By providing reliable high bandwidth access to DOE laboratories and other major research facilities, ESnet4 will enhance the capabilities of researchers and scientists across the country to advance the scientific mission of the Department. Furthermore, because of the infrastructure and design of ESnet4, it will be upgraded seamlessly to meet the growing, complex needs of the Department.

Detailed Justification

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
High Performance Production Computing	**56,700**	**54,790**	**55,000**

This activity supports the NERSC facility located at LBNL. NERSC delivers high-end capacity computing services and supports the SC research community. Nearly 3,100 computational scientists in about 400 projects use NERSC to perform basic scientific research across a wide range of disciplines including astrophysics, chemistry, climate modeling, materials, analysis of data from high energy and nuclear physics experiments, investigations of protein structure, and a host of other scientific endeavors. NERSC enables teams to prepare to use the LCFs as well as to perform the calculations that are required by the missions of the SC programs. NERSC users are supported by Office of Science programs with 60% based in universities, 29% in national laboratories, 6% in other government laboratories, and 3% in industry.

The costs for NERSC fall into three general areas: lease payments for high performance computing hardware, operations (space, power, cooling, maintenance, tapes, etc.), and operating and support staff. NERSCs large user base requires an extremely agile support staff. Careful planning of upgrades is critical to meeting increasing demand within a stable funding profile.

FY 2010 funding will support operation of the NERSC high-end capability systems (NERSC-5 and NERSC-6), lease payments and user support. In FY 2010, the total capacity of NERSC will be approximately one petaflop with the acquisition and operation of NERSC-6. Two systems were decommissioned in FY 2009, Jaquard and Bassi. The Bassi system was replaced with the competitively selected NERSC-6. These computational resources are integrated by a common high performance file storage system that enables users to easily migrate to any of the NERSC resources. With many petabytes of storage and an average transfer rate in the hundreds of megabytes per second this system also allows users to easily move data into and out of the NERSC facility.

	FY 2008	FY 2009	FY 2010
Achieved Operating Hours	8,585	N/A	N/A
Planned Operating Hours	8,585	8,585	8,585
Optimal Hours	8,760	8,760	8,760
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	2,500	3,100	3,100

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
Leadership Computing Facilities (LCF)	**114,065**	**115,000**	**130,000**

The LCF activity was initiated in FY 2004 to enable open scientific applications to harness the potential of leadership computing to advance science. FY 2009 saw the realization of that vision, with four ASCR systems in the top ten of the independent international "Top 500" ranking of supercomputers—including the top system available for open science in two consecutive lists—and the world's first petascale science application which provided new insights into superconductivity. FY 2010 will be the first full year of operation of the petaflop leadership system at Oak Ridge and the 500 teraflop leadership system at Argonne. This new era of petaflop science opens significant opportunities to dramatically advance research as simulations more realistically capture the behavior of materials and ITER scale fusion devices. The success of this effort is built on the gains made in research and evaluation prototypes and the SciDAC program and on years of research in applied mathematics and computer science.

The costs for the LCFs fall into three general areas: lease payments, operations (space, power, cooling, maintenance, tapes, etc), and operating and support staff. In FY 2010, costs are driven by increases in lease payments at both facilities in accordance with the approved baselines.

▪ **Leadership Computing Facility at ORNL (OLCF)**	**84,716**	**85,000**	**88,000**

The first LCF capability for science was established in late FY 2005 at ORNL. In FY 2009, the facility accepted a 1.382 petaflop (1,382 teraflops) Cray Baker system which, when combined with the 263 teraflop machine can achieve a theoretical peak of 1.645 petaflops. The Baker system was accepted in FY 2009 and is expected to complete transition to operations by the end of FY 2009. Some applications, particularly in superconductivity and fusion research, are already achieving sustained performance in excess of 1.3 petaflops or over 80% of the peak capability.

In FY 2010, increases in the lease payments are partially offset by reductions in activities such as enhancing data storage and other site preparation for the petaflop system that were completed in FY 2009. The OLCF activity will support operation and INCITE allocation of the 1.645 petaflop combined system in FY 2010 and will continue to support INCITE projects, ASCR Leadership Computing Challenge projects, scaling tests, and tool and library developers.

	FY 2008	FY 2009	FY 2010
Achieved Operating Hours	7,008	N/A	N/A
Planned Operating Hours	7,008	7,008	7,008
Optimal Hours	7,008	7,008	7,008
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	400	400	400

	(dollars in thousands)		
	FY 2008	FY 2009	FY 2010
▪ **Leadership Computing Facility at ANL (ALCF)**	**29,349**	**30,000**	**42,000**

Diversity in the LCF resources is achieved by operation of the ALCF high performance IBM Blue Gene/P with low-electrical power requirements and a peak capability of 556 teraflops. This facility provides many applications, including molecular dynamics and materials, with access to a system that is better suited to their computing needs than the OLCF and NERSC. In FY 2009, the ALCF completed transition to operation of the IBM Blue Gene/P and provided nine months of operations for INCITE users.

In FY 2010, funding increases support strengthening the Argonne infrastructure and increases in the lease payments on the Blue Gene/P in accordance with the approved schedule. Strengthening the infrastructure in FY 2010, including the start of site preparations, will be essential to prepare the facility for a next generation machine such as the one being developed through the joint research project with NNSA and IBM on the Blue Gene/Q. The ALCF activity will also support operation and INCITE allocation of the Blue Gene/P in FY 2010 and will continue to provide support to INCITE projects, pioneer applications, and tool and library developers.

	FY 2008	FY 2009	FY 2010
Achieved Operating Hours	6,000	N/A	N/A
Planned Operating Hours	6,000	7,008	7,008
Optimal Hours	6,000	7,008	7,008
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	50	100	200

Research and Evaluation Prototypes	**25,347**	**23,900**	**23,900**

The Research and Evaluation Prototype activity addresses the challenges of the systems that will be available by the end of the decade. We anticipate that these systems will be significantly more complex than current computing systems. As a result, many of the tools and techniques developed over the past decade will no longer be effective. By actively participating in the development of these next-generation machines, researchers will better understand their inherent challenges and can begin to work on overcoming those challenges. This activity will prepare researchers to effectively utilize the next generation of scientific computers and will also reduce the risk of future major procurements.

In FY 2010, the Research and Evaluation Prototypes activity will be carried out in close partnership with the NNSA and the Defense Advanced Research Projects Agency (DARPA) program for High Productivity Computing Systems (HPCS). This activity supports completion of both the DOE partnership in the DARPA HPCS Phase III program ($19,500,000) and SC's participation in the joint SC-NNSA partnership with IBM to explore and advance low power density approaches to petascale computing ($4,400,000).

	FY 2008	FY 2009	FY 2010
	(dollars in thousands)		
High Performance Network Facilities and Testbeds	**27,762**	**25,000**	**29,862**

This activity supports operation and upgrades for the Energy Science network (ESnet) including a related research partnership with Internet2. The ESnet provides a high bandwidth network connecting DOE researchers with large-scale scientific user facilities and other scientific instruments. Each year the amount of data generated by these facilities roughly doubles. To meet demand, ESnet has partnered with Internet2—the leading provider of university networks—to push the state-of-the-art and deliver next generation optical network technologies that greatly expand capacity in the core science networks. Continued progress in high performance networks builds on the tools and knowledge developed by the Next Generation Networks for Science research activity and by innovations developed in partnership with Internet2.

The costs for ESnet are dominated by operations which includes refreshing hardware, such as switches and routers, on an accelerated schedule to ensure the 99.99% reliability that is required for large scale scientific data transmission.

In FY 2009, a Recovery Act supported research effort will begin to implement a National testbed of next generation optical technologies that would allow networks, such as ESnet, to use the existing fiber plant but gain a ten fold increase in bandwidth. The testbed will allow ASCR to develop and harden the tools necessary to ensure data integrity and network reliability with this new technology. If successful, the Recovery Act project will allow ASCR to deploy this next generation technology (100 Gbps per wavelength) beginning in FY 2011. This would allow up to 400 Gbps total capacity to many SC laboratories in FY 2011 on a path to achieving 1,000 Gbps connectivity in FY 2014.

	FY 2008	FY 2009	FY 2010
Achieved Operating Hours	8,760	N/A	N/A
Planned Operating Hours	8,760	8,760	8,760
Optimal Hours	8,760	8,760	8,760
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	0.01%	0.01%	0.01%
Number of Users	N/A	N/A	N/A

	FY 2008	FY 2009	FY 2010
Small Business Innovative Research (SBIR)/ Small Business Technology Transfer (STTR)	**—**	**5,925**	**6,446**

In FY 2008, $5,038,000 and $604,000 were transferred to the SBIR and STTR programs respectively. The FY 2009 and FY 2010 amounts shown are the estimated requirements for continuation of the congressionally mandated SBIR and STTR programs.

	FY 2008	FY 2009	FY 2010
Total, High Performance Computing and Network Facilities	**223,874**	**224,615**	**245,208**

Explanation of Funding Changes

	FY 2010 vs. FY 2009 ($000)
High Performance Production Computing	
The increase will support operation of the NERSC high-end capability systems (NERSC-5 and NERSC-6), lease payments, and user support.	+210
Leadership Computing Facilities (LCFs)	
The increase covers increases in lease payments at both facilities in accordance with approved schedule. The increase also supports the start of site preparation at Argonne National Laboratory for acquisition of a next generation machine.	+15,000
High Performance Network Facilities and Testbeds	
The increase will enable ESnet to begin to deliver 100–400 Gbps connections to SC laboratories in FY 2010. The increase in bandwidth is critical to meeting the growing requirements for Department applications and facilities.	+4,862
SBIR/STTR	
Increase in SBIR/STTR due to increase in operating expenses.	+521
Total Funding Change, High Performance Computing and Network Facilities	**+20,593**

Supporting Information

Operating Expenses, Capital Equipment and Construction Summary

(dollars in thousands)

	FY 2008	FY 2009	FY 2010
Operating Expenses	318,274	355,820	394,000
Capital Equipment	23,500	13,000	15,000
Total, Advanced Scientific Computing Research	341,774	368,820	409,000

Funding Summary

(dollars in thousands)

	FY 2008	FY 2009	FY 2010
Research			
National Laboratories	84,537	102,410	107,894
Universities	28,867	32,627	47,119
Other[a]	29,843	38,993	39,125
Total, Research	143,247	174,030	194,138
Scientific User Facility Operations	198,527	194,790	214,862
Total, Advanced Scientific Computing Research	341,774	368,820	409,000

Scientific User Facility Operations

(dollars in thousands)

	FY 2008	FY 2009	FY 2010
NERSC	56,700	54,790	55,000
ESnet	27,762	25,000	29,862
OLCF	84,716	85,000	88,000
ALCF	29,349	30,000	42,000
Total, Scientific User Facility Operations	198,527	194,790	214,862

Facilities Users and Hours

	FY 2008	FY 2009	FY 2010
NERSC			
Achieved Operating Hours	8,585	N/A	N/A
Planned Operating Hours	8,585	8,585	8,585
Optimal Hours	8,585	8,585	8,585
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	2,500	3,100	3,100

[a] Includes $19,500,000 for DOE's partnership in the DARPA HPCS program.

	FY 2008	FY 2009	FY 2010
ESnet			
Achieved Operating Hours	8,760	N/A	N/A
Planned Operating Hours	8,760	8,760	8,760
Optimal Hours	8,760	8,760	8,760
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	0.01%	0.01%	0.01%
Number of Users	N/A	N/A	N/A
OLCF			
Achieved Operating Hours	7,008	N/A	N/A
Planned Operating Hours	7,008	7,008	7,008
Optimal Hours	7,008	7,008	7,008
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	400	400	400
ALCF			
Achieved Operating Hours	6,000	N/A	N/A
Planned Operating Hours	6,000	7,008	7,008
Optimal Hours	6,000	7,008	7,008
Percent of Optimal Hours	100%	100%	100%
Unscheduled Downtime	1%	1%	1%
Number of Users	50	100	200
Total			
Achieved Operating Hours	30,353	N/A	N/A
Planned Operating Hours	30,353	31,361	31,361
Optimal Hours	30,353	31,536	31,536
Percent of Optimal Hours			
Unscheduled Downtime	1%	1%	1%
Number of Users	2,900	3,600	3,700

Scientific Employment

	FY 2008	FY 2009	FY 2010
# University Grants	211	224	260
Average Size	$270,000	$281,000	$296,000
# Laboratory Projects	165	175	180
# Graduate Students (FTEs)	495	533	563
# Permanent Ph.D.s (FTEs)	698	735	766
# Other (FTEs)	239	248	246