Advanced Scientific Computing Research

Funding Profile by Subprogram

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
Advanced Scientific Computing Research		
Mathematical, Computational, and Computer Sciences Research	153,545	174,033
High Performance Computing and Network Facilities	229,654	291,567
Total, Advanced Scientific Computing Research	383,199[a]	465,600

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"
Public Law 111–358, "America COMPETES Act of 2010"

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 (DOE). A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

Background

Imagine exploring the inner workings of a supernova or traveling through time to observe Earth's global climate as it changes. Scientists today can explore these realms thanks to a 100 fold increase in computing power delivered over the past five years and to the software and algorithms developed to harness the power of these forefront computers.

Throughout history, as we have strived to comprehend the mysteries of the universe, mathematics has been an essential tool. It allowed Eratosthenes to determine the circumference of the earth. Newton invented calculus to understand the movement of the planets. Mathematical research in the 1800s laid the groundwork for Einstein's Theory of General Relativity.

Today, advances in mathematics and computing are providing the foundation for models and simulations, which permit scientists to gain new insights into problems ranging from bioenergy and climate change to Alzheimer's disease. ASCR and its predecessor programs have led these advances for the past thirty years by supporting the best applied math and computer science research, delivering world class scientific simulation facilities, and working with discipline scientists to deliver exceptional science.

[a] Total is reduced by $10,801,000, $9,643,000 of which was transferred to the Small Business Innovation Research (SBIR) program, and $1,158,000 of which was transferred to the Small Business Technology Transfer (STTR) program.

"It is generally accepted that computer modeling and simulation offer substantial opportunities for scientific breakthroughs that cannot otherwise—using laboratory experiments, observations, or traditional theoretical investigations—be realized. At many of the research frontiers, computational approaches are essential to continued progress and will play an integral and essential role in much of twenty-first century science and engineering."[a] For the growing number of problems where experiments are impossible, dangerous, or inordinately costly, exascale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data, producing advances in areas of science and technology that are essential to DOE and Office of Science missions and, in the hands of the private sector, drive U.S. competitiveness.

In FY 2009 and FY 2010, ASCR delivered petascale computing power to the open science community, making available two of the world's fastest computers for open-science. More than a dozen DOE applications have achieved sustained performance on these machines in excess of one petaflop revealing new details and nuances in areas such as biofuels, fusion energy, nanoelectronics, fluid dynamics, astrophysics, superconductors, and magnetic materials. These efforts advance critical technologies for the Department and the Nation. For example, the ASCR partnership with Basic Energy Sciences in computational materials have enabled high-density disk drives, new approaches to photovoltaic cells, and understanding of stress corrosion cracking and high-temperature superconductors that could transform our energy infrastructure. ASCR-supported high-performance machines are the culmination of a decade-long effort to build computing architectures with unprecedented speed and capability. At the same time, the development of new mathematical theories and algorithms contributed as much, or more, as the raw computational power of supercomputers to the speed at which difficult scientific problems can be solved computationally. Finally, to ensure that the science and engineering disciplines were ready to use Leadership Computing Facilities, ASCR fostered collaborations among research communities encompassing applied mathematics, computer science, physics, biology, chemistry, and other disciplines.

ASCR supports world-leading basic research in both applied mathematics and computer science focused on areas relevant to high performance computing. The results of this research are brought to the broader scientific community through the Scientific Discovery through Advanced Computing (SciDAC) program. SciDAC, established in 2001, accelerates progress in computational science by breaking down the barriers between disciplines and fostering more dynamic partnerships between applications—such as astrophysics or biology—and the ASCR-supported computer scientists and mathematicians who deeply understand the hardware and software available. Areas of focus for the SciDAC program to date include climate, astrophysics, materials, and fusion energy. These partnerships have been spectacularly successful, with documented improvements in code performance in excess of 10,000 percent.

The success of the SciDAC partnerships within the Office of Science has stimulated the development of new partnerships between ASCR and the Department's technology offices. ASCR-supported researchers are now focusing on the mathematical and computational challenges of the electricity grid and nuclear reactor modeling. In addition, the ASCR applied mathematics program is supporting research in complex natural and engineered systems and uncertainty quantification. This research has relevance to DOE applied programs such as carbon sequestration, wind energy, next generation nuclear reactors, Smart Grid, and fuels from sunlight. Additional partnerships will begin in FY 2012, broadening these efforts.

Science today is by nature increasingly collaborative. Today's collaborations require researchers not only to communicate with each other, but also to exchange large data sets, and often to run complex

[a] "*The Potential Impact of High-End Capability Computing on Four Illustrative Fields of Science and Engineering*" a 2008 report of the National Research Council of the National Academy of Sciences available on line at http://www.nap.edu/catalog.php?record_id=12451

calculations and experiments in locations remote from where the original data are collected or generated. To facilitate the best collaborations for science, ASCR has played a leading role in driving development of the high-bandwidth networks connecting researchers to each other and their data. The invisible glue that binds today's networks across the world together—effortlessly passing billions of searches and trillions of bits from one corner of the earth to another—has roots in ASCR-supported research. ASCR-supported researchers helped establish critical protocols such as TCP/IP on which the current Internet is based. ASCR advanced networking research makes international collaborations such as the Large Hadron Collider and ITER possible and underpins virtual meeting and other commercial collaboration tools.

Looking forward, major changes are underway in the computing industry due to the critical importance of power and energy efficiency in computing systems from the desktop to the exascale. To enable this transition, ASCR is adopting the strategy used by other SC programs and planning more direct involvement with the providers of high performance computing systems. Future computing systems will have thousands of processing units per chip and these units will be heterogeneous. The first generation of this type of high performance hybrid, multi-core computer has already become reality with the emergence of a machine in China that claimed the title of the world's fastest computer in November 2010. The promise of this new hybrid, multi-core architecture brings with it many new challenges. The Chinese machine contains hundreds of thousands of compute cores that share data and communicate with specialized chips that accelerate performance but make it more challenging to use. Soon, hybrid machines will have millions then billons of cores, which need to be managed and made to work effectively together. Communication, data movement, memory management, and fault detection and recovery will all become major challenges for applications. This is not unlike the change from operator run switchboards to computer-run communications except that the users will need to invent and make the change since industry investments to adapt these chips for the gaming and video-streaming markets will not deliver the software, operating systems, or tools needed for scientists to use high performance computers. ASCR must take on that challenge to continue progress in the computational research focused on the Department of Energy missions. However, U.S. scientists also require this capability to retain leadership in many other areas of research. Across computational science and engineering, continued progress will require the expertise of a team of computer scientists, applied mathematicians and computational scientists who deeply understand their needs as well as partnerships with the computer venders to ensure that these needs impact design of new hardware. In so doing, these teams create the software, tools, and methods that advance use of the technology and broaden the impact of its potential throughout the government and private sector.

ASCR will continue to be guided by science needs as it develops computers and networks at the leading edge of technology. ASCR has already initiated investments to address the challenges of hybrid, multi-core computing up to the exascale (capable of an exaflop, or 10^{18} floating point operations per second, a thousand-times faster than today's petascale computers). Like the path to petascale computing, this path is driven by the requirements of applications that are critical to the DOE and the nation as defined through a series of workshops.[a] However, unlike the path to petascale, there are technological challenges that must be addressed to reduce the energy demands and increase the memory available on exascale systems so that they will be useful for science and engineering. Addressing these challenges will result in not only exascale systems but also in affordable, energy efficient petascale systems and high-end desktops to drive scientific and engineering discovery across the country. With this integrated approach, ASCR will continue to deliver scientific insight to address national problems in energy and the environment, while advancing U.S. competitiveness in information technology and high-tech industry.

---

[a] Workshop reports can be found at http://www.science.doe.gov/ascr/ProgramDocuments/ProgDocs.html.

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. 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.
- 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 mathematicians, 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

The evolution of the ASCR program will enable the U.S. to take advantage of the changes in computer hardware technology and deliver computers and networks that are a thousand-fold more energy efficient than today, drive unprecedented improvements in the scientific understanding of areas critical to the future of the Country, and secure a competitiveness advantage in high-tech and information technology industries.

Because computer-based simulation is so important to research programs across SC and the Nation, in addition to its core research program, ASCR invests in partnerships to advance use of high end computing in a wide array of disciplines important to DOE and operates the Leadership Computing Facilities as open user facilities with access determined by merit evaluation of proposals. The next generation of ASCR partnerships will drive unprecedented scientific advances for SC and the Nation. Some examples of applications that rely on both ASCR facilities and research efforts and have important benefits to science and society at large include:

- Computational chemistry and simulation of nanomaterials relevant to energy applications. These applications are funded in partnership with the Basic Energy Sciences program.
- Next generation Earth System Models to dramatically improve our ability to predict changes in global climate. This work is funded in partnership with the Biological and Environmental Research program. 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. This work is jointly funded with the Fusion Energy Sciences program.
- Computer modeling of nuclear structure with relevance for science, nuclear energy, and nuclear weapons. These applications are through partnerships with both the Nuclear Physics program and the National Nuclear Security Administration.
- Analysis of massive amounts of data from experiments such as the Large Hadron Collider, and simulations, such as three dimensional simulations of supernovae events, which are only possible

with leadership computing resources. These works are supported through partnerships with the High Energy and Nuclear Physics programs and the National Nuclear Security Administration.

- Simulations of biological systems relevant for bioenergy applications and subsurface science research to characterize and predict changes in DOE's environmental management sites. This work also has implications for DOE's efforts in subsurface carbon sequestration. These applications are partnerships with the Biological and Environmental Research Program.

In addition to these benefits, establishing SC Leadership Computing Facilities has required partnerships with hardware vendors to develop the most appropriate architectures for scientific discovery and the software necessary to effectively use these powerful systems. These industrial 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.

Finally, through ASCR allocation mechanisms, including INCITE and the ASCR Leadership Computing Challenge (ALCC), industry researchers have used the leadership computing resources to conduct both proof of concept and validation simulations to advance fundamental understanding in their research and development efforts. These users have praised government support for such cutting edge resources and state that their results have helped them gain a competitive advantage by demonstrating the benefits of high performance computing to their companies. Industry applications are held to the same peer review and readiness criteria as academic and national laboratory applications and have come from Boeing, Corning, DreamWorks Animation, Fluent/ANSYS, inc., Ford, Gene Networks Systems, GE Global Research, General Atomics, General Motors Research and Development Center, IBM Research, Pratt and Whitney, Procter & Gamble and Smart Truck/BMI.

Program Planning and Management

ASCR has developed a system of planning and priority setting that strongly benefits from input by 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. One way ASCR ensures the integrity and effectiveness of the peer review processes is through the Advanced Scientific Computing Advisory Committee (ASCAC), which organizes Committees of Visitors (COVs) to review ASCR research management, the impact of ASCR scientific user facilities, and progress toward the long-term goals of the program. For example, a COV reviewed ASCR investments in Applied Mathematics in FY 2010 and found "the program to be very effective and well-managed".[a] In addition, ASCAC identifies scientific challenges and opportunities, including specific bottlenecks to progress in areas such as climate change or computational biology. In FY 2010, ASCAC was charged to analyze the opportunities and challenges for ASCR and the Office of Science associated with exascale computing. In response to this charge, ASCAC formed a subcommittee with experts in a wide range of fields. They considered the challenges and opportunities with primary impact in the Office of Science as well as those impacting DOE more broadly. The subcommittee submitted two interim reports[b] for review and discussion with ASCAC on the path to developing this final report which was unanimously accepted by ASCAC on November 9, 2010. The committee found that "The Exascale Initiative will be significant and transformative for Department of Energy missions and the economic competitiveness of the United States."

In addition to ASCAC, critical tools for managing the ASCR scientific user facilities include annual operational reviews and requirements workshops. For example, ESnet and NERSC both conduct

---

[a] 2010 Applied Math Committee of Visitors report can be found at http://www.sc.doe.gov/ascr/ASCAC/Reports.html

[b] ASCAC *Challenges and Opportunities of Exascale Computing* report can be found at http://www.sc.doe.gov/ascr/ASCAC/Reports.html

requirements workshops each year with individual SC program offices. The purpose of each workshop is to accurately characterize the near-term, medium-term, and long-term requirements of the science conducted by each program office. With two workshops per year, ASCR refreshes the 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, there were a series of workshops in 2009 to identify key science opportunities in the disciplines important to DOE—nuclear energy, materials and chemistry, high energy and nuclear physics, climate change, biology, and cross-cutting areas—and the potential role of extreme scale computing in realizing those opportunities.[a]

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,[b] 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. A recent evaluation of NITRD by the President's Committee of Advisors on Science and Technology (PCAST) recommended balancing procurement of high performance computers with fundamental computer science and engineering research for next-generation technologies. Reports from industry and other mission agencies, such as DARPA, indicate that because of power constraints, data movement, rather than computational operations, will be the limiting factor for future systems petaflops to exaflops. ASCR will use both of these inputs to inform its computer science, and computational science research activities to utilize computing at extreme scales and to understand large scale data from both simulations and experiments, within the context of a mission agency.

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"*[c] 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 high end computing 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. 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."

Basic and Applied R&D Coordination

A cornerstone of the ASCR program is coordination across disciplines and programs. Partnerships within SC are mature and continue to advance the use of high performance computing and scientific

---

[a] Workshop reports can be found at http://www.science.doe.gov/ascr/ProgramDocuments/ProgDocs.html.
[b] Information on the NITRD program can be found at http://www.nitrd.gov.
[c] The "The Potential Impact of High-End Capability Computing (HECC) on Four Illustrative Fields of Science and Engineering" can be found at http://www.nap.edu/catalog.php?record_id=12451

networks for science. In addition, ASCR continues to have a strong partnership with the National Nuclear Security Administration in areas of mutual interest including best practices for management of high performance computing facilities. Looking ahead, this partnership will be increasingly important to efforts to advance the development of exascale computing for critical Department mission applications. In FY 2012, ASCR and NNSA will further strengthen our partnership by coordinating relevant research activities and conducting a series of joint workshops to inform our planning. This coordination will be expanded in FY 2012 to provide for complementary and synergistic efforts among Federal agencies as part of a national strategy for leadership in high performance computing. Through NITRD, ASCR will be a part of a coordinated national effort to address hardware, firmware, and software challenges for exascale computing.

In discussions with the technology development programs throughout DOE, a key area of mutual interest continues to be in applied mathematics for the optimization of complex systems, control theory, and risk assessment. In March 2009, a workshop was organized, in partnership with the DOE Office of Electricity Delivery and Energy Reliability, which focused on the challenges of grid modernization efforts. 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 Offices of Environmental Management, and Fossil Energy), cyber security (with the Office of Electricity Delivery and Energy Reliability), alternative and renewable energy (with the Office of Energy Efficiency and Renewable Energy) and the scientific challenges of exascale computing for national security (with the National Nuclear Security Administration). These workshops facilitate a dialogue between the ASCR research community and a specific applied R&D 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. In FY 2012, ASCR plans to conduct a competition under SciDAC for new applied energy partnerships.

Budget Overview

The architecture of future computing systems, from desktops to exascale, will be transformed by changes in the underlying semiconductor technology and will be constrained by the need for greater energy efficiency. Anticipated characteristics of future hybrid, multi-core systems include:

- Decreased clock frequencies to conserve power;
- Increased number of processing units on a single chip;
- Increased energy costs for moving data both on-chip and off-chip; and
- Decreased memory per flop (and dramatically reduced memory per processor).

Computational scientists will therefore face the following challenges:

- Total concurrency in the applications must rise by a factor of about a million;
- Current weak-scaling approaches will be problematic due to the reduced memory per processor;
- For both power and performance reasons, minimizing data movement will be much more important so flat cache hierarchies will no longer be helpful;
- The failure rates for components and manufacturing variability make it unreasonable to assume the computer is deterministic. This is true for performance today and will affect the results of computations by 2018 due to silent errors;

- Synchronization will be very expensive and the work required to manage synchronization will be high; and
- The I/O system at all levels—chip to memory, memory to I/O node, I/O node to disk—will be much harder to manage due to the relative speeds of the components.

These challenges represent a change in the computing cost model, from expensive compute operations coupled with almost free data movement, to nearly free compute operations coupled with increasingly more expensive data movement.

It is important to note that even though these barriers are the most severe for exascale systems, they impact the full spectrum of computing. Numerical algorithms, mathematical models, and scientific software codes must also be reformulated, perhaps radically, to take full advantage of these emerging computational platforms. New methods of data management and analysis must be developed to deal with the stricter constraints on data movement for visualization and analysis. Achieving the necessary improvements will require a closer cooperation between computer architects, domain scientists, application developers, applied mathematicians, and computer scientists.

The FY 2012 ASCR budget focuses on coordinated efforts to address the fundamental changes taking place in the computing industry to deliver on the promise of hybrid, multi-core computing systems up to the exascale: positioning the U.S. to maintain international scientific computing and simulation leadership to address our scientific and engineering challenges and secure a competitiveness advantage in high-tech and information technology industries. Since the challenges ahead will require major advances in hardware, software, methods, and tools, the request balances investments in high performance computing facilities, advanced networks, and research and evaluation prototypes with investments in applied mathematics, computer science, next generation networks for science, and computational partnerships. This balance should allow for continued progress in a wide array of fields important to DOE's science, energy and national security missions in FY 2012 and for years to come.

The FY 2012 budget request continues support for the Leadership Computing Facility at Oak Ridge National Laboratory (OLCF)—which runs a 2.33 petaflop, six-core Cray XT5 system, openly available to the scientific community through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program as well as the ASCR Leadership Computing Challenge (ALCC) allocations program. 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. In addition, OLCF activities will support preparations at ORNL for a computing system 5-10 times more capable than current systems. These activities are critical to harnessing the complexity of this architecture and meeting challenges associated with hybrid, multi-core high performance computer systems with millions of multicore and specialized processors. OLCF efforts are also expected to build experience and tools for emerging architectures supported in part through ASCR's Research and Evaluation Prototypes activity.

The FY 2012 budget request also continues support for the Leadership Computing Facility at Argonne National Laboratory (ALCF), including an upgrade of this facility to a 10 petaflop IBM Blue Gene/Q. Development of this proposed machine is based on a joint ASCR-NNSA sponsored research project with IBM and the Argonne and Lawrence Livermore National Laboratories, supported from FY 2006 to FY 2010 through the Research and Evaluations Prototypes activity. This project focused on the development of power efficient high performance computing and meaningfully contributed to the IBM Blue Gene product line, which was recognized with the National Medal of Technology and Innovation in the Fall of 2009. The Blue Gene Q provides an alternative approach to high performance computing from the hybrid, multi-core approach being pursued at the OLCF. This diversity both broadens the user community to a wide range of applications important to the Department and also reduces the risk of future procurements. The ALCF is also openly available to the scientific community through INCITE and ALCC and will continue to provide access and assistance to tool and library developers and to researchers seeking to scale their applications. The National Energy Research Scientific Computing Center (NERSC) facility at Lawrence Berkeley National Laboratory will continue to operate at a capacity of over one petaflop in FY 2012 to meet ever growing demand from SC researchers. The focus will be on assisting researchers to effectively utilize the potential of this facility and preparing researchers to move beyond NERSC to the leadership computing machines.

The FY 2012 budget request supports operation of ESnet to continue to advance the next generation network capability that is critical to DOE applications and facilities. Building on the Recovery Act-supported Advanced Networking Initiative, ESnet will begin to deliver 100 gigabit per second (Gbps) connections to SC laboratories in FY 2012, with a goal of achieving 1,000 Gbps connectivity in 2016. The increases in bandwidth are timed to meet the requirements of the Office of Science 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 volume of data in climate research, nuclear structure, genomics, and proteomics that advance DOE's energy and environment missions. These efforts also rely upon investments in Next Generation Networking for Science to ensure that facilities are able to fully utilize this expanded bandwidth.

The FY 2012 budget greatly expands research efforts under Computational Partnerships aimed at developing applications capable of utilizing high performance hybrid, multi-core computing resources while informing the development of exascale resources to ensure that the demands of these applications are well understood and helping to shape the new architectures under development. New Co-Design Centers will be supported in key applications critical to Department missions in energy, environment and national security. These Co-Design Centers will enable close coupling of applications, computer science and computer hardware architecture that is required for success at exascale. In addition, FY 2012 begins the third round of investments in Scientific Discovery through Advanced Computing (SciDAC). To make it easier for science and engineering applications to utilize and incorporate enabling technologies, the 14 assorted SciDAC Centers and Institutes, supported in FY 2010, will be replaced in FY 2012 with a few larger, more integrated, SciDAC Institutes that will develop and deploy the tools and capabilities needed by both the traditional Scientific Computation Application Partnerships, funded with the other SC programs, and expanded Department of Energy Computation Application Partnerships, funded with the Department's applied programs, that enable scientists to refine and scale their codes to effectively utilize the capabilities of the Leadership Computing Facilities.

The FY 2012 budget continues the core research efforts in Applied Mathematics and Computer Science focused on long-term research necessary for continued progress in high performance computational science and engineering. These investments will be critical to meeting the challenges of high performance hybrid, multi-core computing systems and to realizing the potential of the exascale systems on the horizon. The Applied Mathematics activity will support new efforts in Uncertainty Quantification to enable predictive simulations. In networking, the focus will continue to be on developing the advanced tools that harness new optical technologies to meet the requirements of the ESnet and the SC research community.

Significant Program Shifts

The challenges of developing and utilizing high performance hybrid, multi-core computing systems on the path to exascale while meeting the demands of some critical DOE applications require us to enhance the long-term, coordinated research and hardware investments across the ASCR research portfolio that were initiated in FY 2010. The FY 2012 budget enhances these investments with significant new projects in the Research and Evaluations Prototypes activity that focus on understanding and influencing critical technology developments in areas such as memory and processor design to provide increased energy efficiency and performance. These projects also provide the U.S. research community with an opportunity to experiment with cutting-edge architectures and begin to develop tools and methods for harnessing their capabilities. The changes in computer architectures discussed above require increased innovation and collaboration for continued progress. By actively participating in the development of these next-generation machines, ASCR can ensure that appropriate architectures for science are developed while researchers 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. These investments are also important to continued U.S. leadership in information technology and high-tech industry as vendor partnerships ensure that the solutions developed for scientific computing are also incorporated into commercial offerings.

In FY 2012, SciDAC partnerships with DOE's applied technology offices will be expanded. The Computational Partnerships activity will also expand support for Co-Design Centers. The Applied Mathematics activity will support new efforts in Uncertainty Quantification to enable predictive simulations. The Argonne Leadership Computing Facility will be upgraded to approximately 10 petaflops. The Oak Ridge Leadership Computing Facility will prepare for a computing system 5-10 times more capable than the current Cray XT5.

Annual Performance Targets and Results

The Department is in the process of updating its strategic plan, and has been actively engaging stakeholders including Congress. The draft strategic plan is being released for public comment concurrent with this budget submission, with the expectation of official publication this spring. The draft plan and FY 2012 budget are consistent and aligned. Updated measures will be released at a later date and available at the following link http://www.mbe.doe.gov/budget/12budget.

Mathematical, Computational, and Computer Sciences Research

Funding Schedule by Activity

(dollars in thousands

	FY 2010 Current Appropriation	FY 2012 Request
Mathematical, Computational, and Computer Sciences Research		
Applied Mathematics	43,698	48,973
Computer Science	45,936	47,400
Computational Partnerships	49,538	60,036
Next Generation Networking for Science	14,373	12,751
SBIR/STTR	0	4,873
Total, Mathematical, Computational, and Computer Sciences Research	153,545	174,033

Description

The Mathematical, Computational, and Computer Sciences Research subprogram supports activities aimed at effectively utilizing forefront computational and networking capabilities. Computational science is increasingly central to progress at the frontiers of science and to our most challenging engineering problems. Accordingly, the subprogram must be positioned to address scientifically challenging questions, such as:

- What new mathematics are required to more accurately model systems involving processes taking place on vastly different time and length scales such as the earth's climate and the behavior of living cells?
- Which computational architectures and platforms will deliver the most benefit for the science of today and the future?
- What computer science and algorithm innovations would increase the efficiency of supercomputer problem solving?
- What operating systems, data management, analyses, representation model development, user interfaces, and other tools are required to make effective use of future-generation supercomputers?
- What networking collaboration tools are needed to make all scientific resources readily available to scientists, regardless of whether they are in a university, national laboratory, or industrial setting?

FY 2010 Accomplishments

- *Mathematics for Reliable Real-Time Power Dispatch.* Since electricity cannot currently be effectively stored on a large scale, providers rely on usage forecasts to plan generating capacity. Mathematicians and computer scientists at Argonne National Laboratory have recently developed a program called PIPS for the electricity dispatch problem. PIPS combines information such as weather, loads, and renewable generation outputs (e.g., wind power) over wide geographical regions to enable electricity usage forecasts to ensure sufficient capacity at any given moment.

Preliminary numerical testing has demonstrated a 600-fold improvement in the time-to-solution for a dispatch problem with 27 million variables and 6,000 uncertain scenarios using 1,000 parallel cores.

- *New Programming Model for Achieving Scalability on Extreme-Scale Computing Systems.* A critical challenge for future high-performance computing systems is the effective management of the speed and expediency of communications among the extremely large number of processors. ASCR researchers have developed an approach that holds promise for circumventing message passing bottlenecks. Called the Asynchronous Dynamic Load Balancing Library (ADLB), this innovative programming model enables an application to reuse most of its highly tuned code and streamlines the way parallel programming resources are allocated. Although more research will be needed to fully explore the capability of ADLB, early results are promising. Using ADLB, the researchers have successfully run the first, full calculations of carbon-12—a specific target of the ASCR SciDAC Universal Nuclear Energy Density Functional project and one that addresses fundamental issues in the origin of the universe—on more than 130,000 processors of the Blue Gene/P at the Argonne Leadership Computing Facility.
- *New Algorithm Improves Performance and Accuracy on Extreme-Scale Computing Systems.* On modern computer architectures, communication between processors takes longer than the performance of a floating point arithmetic operation by a given processor. ASCR researchers have developed a new method, derived from commonly used linear algebra methods, to minimize communications between processors and the memory hierarchy, by reformulating the communication patterns specified within the algorithm. This method has been implemented in the TRILINOS framework, a highly-regarded suite of software, which provides functionality for researchers around the world to solve large scale, complex multi-physics problems.
- *Improving Community Code Performance for Nuclear Energy Application.* Nuclear reactor analysis requires accurate characterization of the neutron distribution in the reactor in order to determine power, safety, and fuel and component performance. In a steady-state operational reactor, the neutron field is characterized by six independent variables (three in space, two in angle, and one in energy), and the mean free path of low-energy neutrons are in the millimeter to centimeter range. Thus, high-resolution solutions of the transport equation require tremendous computational resources. Traditionally, computational resources have not been sufficient to attack this problem at full resolution; so multi-level approximation schemes have been employed. Denovo is a community code that serves as a general radiation transport application for nuclear and radiological sciences by finding accurate numerical solutions to the linear Boltzmann equation. This application area includes, but is not limited to, nuclear reactor analysis, fusion, radiation shielding and protection, nuclear safeguards, radiation detection, and radiation therapy, diagnostics, and treatment planning. Working with Denovo users, ASCR researchers made major improvements to traditionally used methods for decomposition and for solving the partial differential equations governing the power distribution in a full-scale reactor core. The core (height 4.2 m) has 17x17 (289) assemblies (height 3.6 m), of which 157 are fuel and 132 are reflector, with each fuel assembly consisting of 17x17 fuel pins. As a result, accuracy was dramatically improved by eliminating multi-level approximation schemes and the problem was solved in 11 minutes on the OLCF using over 50% of available computing power. The time to solution was a factor of 17 improvement over previous simulations.

- *Simulation Reveals Details in Cellulose for Bioenergy Alternatives.* ASCR researchers recently conducted cellulose simulations showing a structure change in cellulose when heated that may impact how hard it is to break cellulose apart. ASCR researchers also developed new methods to measure the intrinsic work for an enzyme to pull a single cellulose chain out of a crystal of cellulose—a key measure of the intrinsic toughness of cellulose. Cellulose is the most common organic compound on Earth and is a primary target for bioenergy. Nature has evolved to protect plants from invaders by making cellulose very difficult for enzymes to deconstruct. Individual cellulose chains pack together into fibrils, similar to how small threads are combined to form a rope or cable. Experiments are limited in what they can tell us about this structure, so we rely on molecular simulation to investigate the detailed structure of cellulose fibrils. The simulation results will allow us to quantify the effects of pretreatment strategies on cellulose structure with significant implications for advancing development of cellulosic bio-ethanol.

Detailed Justification

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
Applied Mathematics	**43,698**	**48,973**

Over the past half-century, ASCR-supported researchers from universities, national laboratories, and industry have made significant enduring advances in applied mathematics, such as laying the foundations for Computational Fluid Dynamics and developing Adaptive Mesh Refinement for tracking important, fine-scale structures within a simulation. These have been essential enablers of modern computational science and engineering and have become ubiquitous in industry and throughout the government. The challenges of the Department's missions and those from high performance hybrid, multi-core computing will require significant innovation from our applied mathematicians and a dedicated focus throughout our portfolio. These anticipated results will transform computational and applied mathematics and will be critical to continued leadership in computational science. As the computing industry moves toward hybrid, multi-core computing architectures, the applied mathematics-supported researchers will be working with the software industry to continue progress at all levels of computing with implications throughout our economy and our daily lives.

The Applied Mathematics activity supports the research and development of applied mathematical models, methods, and algorithms for understanding complex natural and engineered systems related to DOE's mission. These mathematical models, methods, and algorithms are the fundamental building blocks for describing physical and biological systems computationally. Applied Mathematics research underpins all of DOE's modeling and simulation efforts. The research falls into several general categories:

- Numerical methods research for equations related to problems such as fluid flow, magneto-hydrodynamics, wave propagation, and other natural or physical processes.
- Computational meshing research for developing ways in which physical domains can be efficiently partitioned into smaller, possibly geometrically complex, regions as part of a larger-scale simulation.

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request

- 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.
- Optimization research for mathematical methods for minimizing energy or cost, 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.
- In FY 2010, the Multiscale Mathematics area of research began to explicitly include Uncertainty Quantification. For FY 2012, this area is expanded to a somewhat separate research area owing to its importance for predictive science at the exascale.
- 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.

In FY 2012, basic research activities will continue for fundamental mathematical advances and computational breakthroughs across DOE and Office of Science missions—including model formulation and algorithm development for emerging computing systems. FY 2012 supports significant growth in support for Uncertainty Quantification research, deemed critical to realizing the potential of exascale computing for predictive science. For many problems in the fields of natural sciences and engineering, incomplete descriptions, measurements, and data for these problems introduce uncertainties as researchers attempt to understand complex phenomena through computer modeling and simulation. As we increase the use of computing to study such problems, the development of more sophisticated techniques for the incorporation and quantification of uncertainties becomes key to our success. New efforts in the research and development of uncertainty quantification methodology and techniques will improve our overall understanding of complex scientific and engineering problems and allow us to make quantitative predictions about the behavior of these systems.

In FY 2012, the Computational Science Graduate Fellowship program, aimed at attracting the best graduate students in the scientific disciplines and educating them as the next generation of computational scientists, is continued at $6,000,000.

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request
Computer Science	**45,936**	**47,400**

Computer Science research is conducted by universities, national laboratories, major information technology companies and small businesses. Their efforts have resulted in software, languages, libraries and tools that run on most of the world's high performance computing platforms and are used by many of the Nation's information technology and high-tech companies, as well as by all of the DOE scientific and national security application communities. As the computing industry moves toward hybrid, multi-core computing architectures, this work plays a critical role in ensuring that these machines are useful tools for all levels of scientific computing.

The Computer Science activity supports research to utilize computing at extreme scales and to understand extreme scale data from both simulations and experiments. Industry reports indicate that because of power constraints, data movement, rather than computational operations, will be the limiting factor for future systems petaflops to exaflops. Memory per core is expected to decline sharply while the performance of storage systems will continue to lag far behind the computational capability of the systems. Multi-level storage architectures that span multiple types of hardware are anticipated and will require new approaches to run-time data management and analysis. To address these challenges the Computer Science portfolio includes:

- 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.
- Approaches to simulate and understand the impact of advanced computer architectures on scientific applications critical to the Department.
- 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.

A fundamental challenge for the Computer Science activity is enabling science applications to harness computer systems with increasing scale and increasing complexity that take advantage of technology advances such as multicore chips. This challenge will require more dynamic behavior of system software (operating systems, file systems, compilers, and performance tools) than historically developed to deal with time varying power and resilience requirements. Substantial innovation is

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request

needed to provide essential system software functionality in a timeframe consistent with the anticipated availability of hardware.

Another fundamental challenge is enabling scientists to effectively manage, analyze, and visualize the petabytes of data resulting 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 2012, the Computer Science activity will continue to focus on the challenges of emerging high performance hybrid, multi-core computing architectures containing as many as a billion cores with hybrid processors. The research efforts, begun in FY 2010, will continue to focus on advanced architectures and related technologies for hybrid, multi-core computing up to the exascale and the associated software including significant investments to improve our ability to simulate future systems.

Computational Partnerships	**49,538**	**60,036**

Over the past decade, the Scientific Discovery through Advanced Computing (SciDAC) has influenced and shaped the development of a distinct approach to science and engineering research through high performance computation. The influence of this program can best be seen by the world-wide interest in replicating the SciDAC approach. SciDAC focuses on the very high-end of high performance computational science and engineering and faces two distinct challenges: to broaden the community and thus the impact of this approach, particularly to address the Department’s missions, and to ensure that further progress at the forefront is enhanced rather than curtailed by the emergence of hybrid, multi-core architectures. A decade of effort has uniquely enabled this program to be ready to simultaneously meet both of these important challenges. These efforts have shown U.S. industry new ways to use their computing to improve competitiveness. The Computational Partnerships activity supports the SciDAC program to use results from applied mathematics and computer science research on scientific applications sponsored by other SC programs. These partnerships enable scientists to conduct complex scientific and engineering computations on leadership-class and high-end computing systems at a level of fidelity needed to simulate real-world conditions. To address challenges of the future, SciDAC has been expanded to include Co-Design Centers/partnerships focused on developing the applications that need exascale computing systems and informing the design of the hardware.

The next round of SciDAC investments, scheduled for competition in FY 2012, will address the challenges from emerging hardware to ensure continued progress in computational science in support of the Department’s missions and broadening the benefits of high end computation to the Department’s applied programs. The key components of the program are SciDAC Institutes, Computation Application Partnerships, and Co-Design Centers.

- *SciDAC Institutes* will be the keystone for applied math and computer science efforts to systematically address technical challenges that are inherent to the scale of new architectures or common across applications. The formulation of these Institutes is based on our lessons-learned from nearly a decade of SciDAC Centers for Enabling Technology and University-based Institutes. The re-formulated SciDAC Institutes will be responsible for developing 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 tools for feature identification, data management, and visualization spanning the full range of SciDAC applications. The SciDAC Institutes will provide a more integrated approach for managing outreach to new communities and deployment of SciDAC-developed software and techniques.

- *Scientific Computation Application Partnerships* with other SC programs will build on past successes to dramatically improve the ability of SC researchers to effectively and confidently utilize multi-petaflop computing systems to advance science. These partnerships support research between applied mathematicians and computer scientists (supported by ASCR) with domain scientists (supported by the other SC 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.
- *Department of Energy Computation Application Partnerships* with the Department's applied programs will be formed to transfer SciDAC technologies and improve the ability of DOE researchers to effectively utilize petascale and emerging extreme scale computing systems to advance their mission goals. These partnerships support collaborative research between applied mathematicians and computer scientists (supported by ASCR) with application researchers (supported by the other DOE programs) to refine and apply computational techniques and tools that address the specific problems of a particular research effort, such as modeling the complete nuclear fuel cycle or the electricity grid; developing a predictive model that accurately simulates conditions at DOE contaminated sites and allows the Department to optimize remediation strategies; and predictive, long-term simulations of carbon sequestration options. ASCR has had workshops with NE to inform these partnerships and is currently planning workshops with EERE, OE, and FE to develop these efforts.
- *Co-Design Centers* will continue the close coupling of applications, computer science, and computer hardware architecture that is required for success at exascale. They ensure that future architectures are well-suited for DOE target applications and that major DOE scientific problems can take advantage of the emerging hybrid, multi-core computing architectures. These centers of computational science need to be formally engaged in the hardware, software, numeric methods, algorithms, and applications co-design process that will be responsible for making key tradeoffs in the design of exascale systems. It is anticipated that these Co-Design Centers will need to address: application formulation, advanced programming languages, integrated uncertainty quantification, validation and verification, new mathematics and approaches to implementing multi-physics problems that naturally express parallelism and locality, and data management analysis, including visualization, all in the context of a scientific problem area important to the DOE.

In FY 2012, the request supports new Co-Design efforts and new Department of Energy Computation Application Partnerships with the Department's applied programs for energy and the environment.

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request
Next Generation Networking for Science	**14,373**	**12,751**

To facilitate the best collaborations for science, ASCR has played a leading role in driving development of the high-bandwidth networks connecting researchers to each other and their data. The invisible glue that binds today's networks across the world together—effortlessly passing billions of searches and trillions of bits from one corner of the earth to another—has roots in ASCR-supported research, ASCR-supported researchers helped establish critical protocols such as TCP/IP on which the current Internet is based. ASCR advanced networking research makes international collaborations such as the Large Hadron Collider and ITER possible and underpins virtual meeting and other commercial collaboration tools. 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:

- 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 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; and
- 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 2012, research will continue to focus on developing the software, middleware and hardware that delivers 99.999% reliability up to and beyond one petabyte per second while allowing the successful products of prior research such as the Open Science Grid and the Earth Systems Grid to transition into operational services supported by their users. The Next Generation Networking research program will continue to make critical investments in several areas including (1) new protocols that allow hosts to rapidly and efficiently adapt to network conditions to maximize the available bandwidth; (2) new routing algorithms that can improve the performance of routers and switches; (3) a rich suite of secure collaboration tools and services; and (4) advanced simulation environments that duplicate real networks to ensure that science communities achieve their goals, and in so doing, underpin commercial offerings to broaden the opportunities for other government agencies, U.S. industry and the American people.

Small Business Innovative Research (SBIR)/ Small Business Technology Transfer (STTR)	**0**	**4,873**

In FY 2010, $4,113,000 and $494,000 were transferred to the SBIR and STTR programs respectively. The FY 2012 amount shown is the estimated requirement for continuation of the SBIR and STTR programs.

Total, Mathematical, Computational, and Computer Sciences Research	**153,545**	**174,033**

Explanation of Funding Changes

	FY 2012 vs. FY 2010 Current Approp. ($000)
Applied Mathematics	
This increase will support new efforts in Uncertainty Quantification to enable predictive scientific simulation and the continued development of new and improved applied mathematical models, methods, and algorithms to understand complex systems.	+5,275
Computer Science	
The increase will support new efforts focused on the challenges of hybrid, multi-core computing.	+1,464
Computational Partnerships	
The increase, in addition to the completion of several projects, will support new partnerships with DOE applied programs and Co-Design Centers.	+10,498
Next Generation Networking for Science	
This decrease reflects a phase out of support for the Open Science Grid and Earth Systems Grid which are moving into production.	-1,622
SBIR/STTR	
SBIR/STTR increases as research funding is increased.	+4,873
Total Funding Change, Mathematical, Computational, and Computer Sciences Research	**+20,488**

High Performance Computing and Network Facilities

Funding Schedule by Activity

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request
High Performance Computing and Network Facilities		
High Performance Production Computing	54,900	57,800
Leadership Computing Facilities	128,788	156,000
Research and Evaluation Prototypes	15,984	35,803
High Performance Network Facilities and Testbeds	29,982	34,500
SBIR/STTR	0	7,464
Total, High Performance Computing and Network Facilities	229,654	291,567

Description

The High Performance Computing and Network Facilities subprogram delivers forefront computational and networking capabilities to scientists nationwide. These include high performance production computing at the National Energy Research Scientific Computing Center (NERSC) facility at LBNL and Leadership Computing Facilities at Oak Ridge (OLCF) and Argonne (ALCF) National Laboratories. These computers, and the other SC research facilities, generate many petabytes of data each year. Moving data to the researchers who need them requires advanced scientific networks and related technologies provided through High Performance Network Facilities and Testbeds which includes the Energy Science network (ESnet). The subprogram also invests in long-term needs through the Research and Evaluation Prototypes activity which will play a critical role in achieving exascale computing.

Computing resources are allocated through competitive processes. Up to eighty percent of the processor time on the Leadership Computing Facilities is allocated through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, which is open to all researchers and results in awards to 20–30 projects per year, each requiring a substantial amount of the available resources. The high performance production computing facilities at NERSC are predominately allocated to researchers supported by SC programs. In addition to INCITE, all of the ASCR scientific computing facilities allocated ten to thirty percent of computing resources through the ASCR Leadership Computing Challenge program. This program is open year-round to scientists from the research community in academia and industry for special situations of interest to DOE with an emphasis on high-risk, high-payoff simulations in areas directly related to the DOE's energy mission, for national emergencies, or for broadening the community of researchers capable of using leadership computing resources.

ASCR facilities provide critical resources for the scientific community following the public access model used by other Office of Science scientific user facilities including light sources, accelerators and microscopes. In addition, ASCR facilities provide a crucial testbed for U.S. industry to deploy the most advanced hardware and have it tested by the leading scientists across the country in universities, national laboratories, and industry.

FY 2010 Accomplishments

- *Oak Ridge Supercomputers Delivers Petascale Science*. Using the additional computing capability from the FY 2009 upgrade to 2.33 petaflops, researchers using the Oak Ridge Leadership Computing Facility (OLCF) were able for the first time to predict the flux of contaminants into the Columbia River basin at the Hanford 300 waste disposal site in Washington State. The team's application was able to model the complex nonlinear chemistry and transient fluid flow caused by fluctuations in the Columbia River to inform the Department's efforts for the remediation of the aging underground contaminant storage facilities. Another project run on the OLCF looked at lifted flames from a direct injection diesel-engine jet and, for the first time, performed a three-dimensional simulation capable of fully resolving flame and ignition features such as chemical composition, temperature profile, and turbulence flow. In yet another project, researchers performed the largest simulation to date of turbulent dispersion, informing the development of more efficient engines.
- *Argonne Supercomputer Explains How Impurities Cause Embrittlement of Materials.* Simulations run at the Argonne Leadership Computing Facility have answered a fundamental question encompassing chemistry, mechanics, and materials science: how minute impurities in the grain boundaries of a material alter its fracture behavior. Sulfur segregation-induced embrittlement of nickel is an important problem for the design of the next-generation nuclear power plants. Although experiments had identified an essential role of sulfur in the embrittlement, there was little understanding of how the sulfur had effect. Through simulation, researchers found that the impurities led to an order-of-magnitude reduction of grain-boundary shear strength and tensile-strength reduction, which allow crack tips to find an easy propagation path. This mechanism explains experimental observations and elucidates the experimentally found link between sulfur impurities and embrittlement. Results were published in *Physics Review Letters*.
- *Argonne Supercomputer Speeds up Research in Response to Pandemic Flu.* The rapid turnaround of huge protein structure campaigns at the Argonne Leadership Computing Facility allowed researchers seeking to computationally design protein-based inhibitors towards pathogens like the H1N1 influenza to develop and test new algorithms that speed up research. The acceleration provided by use of leadership computing resources allowed researchers to quickly identify two candidate inhibitors for the pandemic strain of H1N1 in response to the global emergency. Conventional protein structure determination from nuclear magnetic resonance (NMR) data is labor intensive and prone to error. Researchers used the ALCF to develop simplified approaches without losing accuracy. These approaches were incorporated into the Rosetta3 protein structure modeling methodology—the most popular software tool for protein structure prediction. This enhancement enables routine NMR structure determination for larger proteins. Results were published in *Science* and the *Journal of Molecular Biology*.
- *National Energy Research Scientific Computing (NERSC) Clears Cloud Modeling*. Researchers, using supercomputers at NERSC, developed a new global cloud resolving model which is designed to reduce the uncertainty in current climate models. Combined with these powerful systems of the future, the model will allow scientists to simulate climate down to the one-kilometer scale—the scale at which the effects of clouds will be more accurately represented. Clouds have powerful effects on the Earth's climate, both cooling the planet by reflecting heat from the sun and warming it by trapping heat near the surface. Today's models, at best, use parameterizations, or statistical representations, of clouds and their effects in climate models—one of the main reasons different

climate models produce different results. In fact, cloud parameterizations are one of the greatest sources of uncertainty in today's climate models.

- *Energy Sciences Network (ESnet) Delivers*. When the Large Hadron Collider started operations in FY 2010, ESnet flawlessly handled the doubled traffic over two months, transmitting over 8.6 petabytes of data per month. More than half of this traffic is now using the software developed by ESnet to establish virtual circuits that provide carrier-class network reliability and services. ESnet has also deployed a test and measurement infrastructure that significantly increased data transfer performance for the climate community, DOE supercomputer centers and leadership facilities in support of the Intergovernmental Panel on Climate Change (IPCC) 5th assessment report (AR5) expected in 2014.

Detailed Justification

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
High Performance Production Computing	**54,900**	**57,800**

This activity supports the NERSC facility located at LBNL. NERSC delivers high-end capacity computing services for the SC research community. Annually, nearly 3,200 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, high energy and nuclear physics, and biology. NERSC enables teams to prepare to use the ALCF and OLCF as well as to perform the calculations that are required by the missions of the SC programs. NERSC users are supported by SC programs with about 60% based in universities, 30% in national laboratories, and 10% in other government laboratories and industry. NERSC's large and diverse user base requires an extremely agile support staff.

FY 2012 funding will support operation of the NERSC high-end capability systems (NERSC-5 and NERSC-6), lease payments, and user support. The total capacity of NERSC in FY 2012 will be more one petaflop. The NERSC computational resources are integrated by a common high performance file storage system that enables users to easily migrate to any of the available resources. With approximately 60 petabytes of storage and an average transfer rate in the hundreds of megabytes per second, this system is critical to effective use of the facility and allows users to easily move data into and out of the NERSC.

The NERSC facility is a vital resource for the Office of Science research community and is consistently oversubscribed with requests exceeding capacity by a factor of 3-10. NERSC regularly gathers requirements from the Office of Science programs through a robust process that informs upgrade plans. Regular upgrades attempt to keep pace with demand but scientific demand for computing continues to outpace available resources. NERSC is also a vital resource for ASCR because it has a large and diverse user community. NERSC staff are adept at meeting the needs of such a diverse community including both savvy and novice users. This feedback is vital to SciDAC and to ASCR efforts to broaden the impact of its research results particularly as the computing industry moves toward hybrid, multi-core computing.

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request

	FY 2010	FY 2012
Achieved Operating Hours	8,585	N/A
Planned Operating Hours	8,585	8,585
Optimal Hours	8,585	8,585
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	3,100	3,500

Leadership Computing Facilities	**128,788**	**156,000**

The Leadership Computing Facilities (LCF) activity enables open scientific applications to harness the potential of leadership computing to advance science. In June, 2010, the Oak Ridge LCF was the most powerful computing system in the world at that time, according to the independent international Top 500 ranking of supercomputers. The era of petaflop science opens significant opportunities to dramatically advance research as simulations more realistically capture the behavior of, for example, materials and ITER scale fusion devices. The leadership facilities also provide experience and tools critical to continued progress in high performance computing. The success of this effort is built on the gains made in research and evaluation prototypes, the SciDAC program, and research in applied mathematics and computer science.

The costs for both the ALCF and OLCF fall into three general areas: lease payments, operations (space, power, cooling, maintenance, tapes, etc.), and staff.

▪ **Leadership Computing Facility at ORNL (OLCF)**	**86,788**	**94,000**

In FY 2009, with Recovery Act funds, the OLCF’s Cray XT5 system was upgraded to 2.33 petaflops—making it the most powerful computer in the world. In addition, the facility also continues to operate a 263 teraflop Cray XT4 machine, also available for INCITE projects, ASCR Leadership Computing Challenge projects, scaling tests, and tool and library developers.

The facility staff continues to assist users to fully utilize the OLCF resources. As a result, several applications, such as combustion studies in diesel jet flame stabilization, simulations of neutron transport in fast reactor cores, and groundwater flow in porous media, are running at petascale.

In FY 2012, the request supports staff, operations, and lease payments for the computer system at the facility. In addition, the OLCF will continue site preparations for a 5-10 times more capable computing system. Experience with this system is expected to inform ASCR investments in both exascale relevant research and in the Research and Evaluations Prototypes activity.

The OLCF has a strong and successful partnership with its computing vendor. The OLCF’s knowledgeable staff are versed in hardware, cutting edge tools and methods, and the needs of forefront computational science and engineering teams. It is this expertise that enabled the OLCF to deliver petascale science from day one of operations of the petascale computer. The OLCF then turned its focus on the specific challenges of hybrid, multi-core computing systems with particular attention to application requirements. Through INCITE and outreach, the OLCF transfers its expertise to industry and academic users to broaden the benefits for the Nation. For example, the OLCF worked with Ford on energy recovery technologies and is working with BMI/Smart Truck to improve the fuel efficiency of class 8 trucks.

	FY 2010	FY 2012
Achieved Operating Hours	7,008	N/A
Planned Operating Hours	7,008	7,008
Optimal Hours	7,008	7,008
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	550	625

- **Leadership Computing Facility at ANL (ALCF)** **42,000** **62,000**

The ALCF provides a high performance 556 teraflop peak capability IBM Blue Gene/P with low-electrical power requirements. It will be upgraded in FY 2012 to the next generation system, an IBM Blue Gene/Q, with peak capability of approximately 10 petaflops. The Blue Gene/Q was developed through a joint research project with NNSA, IBM, and ASCR's Research and Evaluation Prototypes activity. The ALCF and OLCF systems are architecturally distinct and this diversity of resources benefits the Nation's HPC user community. The ALCF supports many applications, including molecular dynamics and materials, for which it is better suited than the OLCF or NERSC. ALCF staff operates and maintains the computing resources and provides support to INCITE projects, ASCR Leadership Computing Challenge projects, scaling tests, and tool and library developers.

In FY 2012, the request supports acquisition, installation and testing of the IBM Blue Gene/Q. In addition, the ALCF activity will also support operation and INCITE allocation of the Blue Gene/P in FY 2012 and operation and early access to the Blue Gene/Q.

The ALCF also has a strong and successful partnership with its computing vendor. The ALCF's knowledgeable staff are versed in hardware, cutting edge tools and methods, and the needs of forefront computational science and engineering teams. When IBM received the National Medal of Technology and Innovation for its Blue Gene line of supercomputers in the Fall of 2009, the contributions of Argonne and Lawrence Livermore National Laboratory, from their research partnership with IBM, were also recognized. This experience will be critical to the success of industry partnerships to address the challenges of next-generation computing. Through INCITE, the ALCF also transfers its expertise to industry and academic users to broaden the benefits for the Nation. For example, the ALCF has worked with Pratt and Whitney to reduce emissions and improve operability in aircraft engine combustors and it is working with GE Global Research to reduce Aerodynamic noise in next-generation, low-emission aircraft propulsion and wind turbines.

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request

	FY 2010	FY 2012
Achieved Operating Hours	7,008	N/A
Planned Operating Hours	7,008	7,008
Optimal Hours	7,008	7,008
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	200	350

Research and Evaluation Prototypes	**15,984**	**35,803**

The potential benefits from exascale computing capability have been documented through a series of workshops and town meetings. These workshops and detailed discussions with industry experts underscore that achieving exascale computing will require overcoming large technical challenges that cannot be accomplished with incremental improvements to current hardware and software. One of the most significant challenges will be reducing the amount of electrical power per operation by approximately 99.7% relative to today's systems. Memory design is a related challenge as future systems must dramatically reduce the amount of energy needed to move data between the memory devices and the processor. Another hardware challenge arises from the continued reduction of transistor dimensions which will cause many circuit failures. Techniques will be needed to manage the errors and it may be necessary for programming techniques and operating system designs to handle those errors that are behind hardware detection and correction mechanisms. Additional hardware challenges include the move to solid-state disk storage, the use of on-die voltage regulation, and the transition to photonic interconnection networks for moving the massive amount of information between processor-memory elements and between those elements and bulk storage and networking devices. All in all, the hardware challenges of exascale dwarf those faced in moving from gigascale to terascale in the 1990s and in moving from terascale to petascale in the last decade.

The Research and Evaluation Prototypes activity addresses these challenges of next generation computing systems. By actively partnering with industry in the development of these next-generation machines, ASCR can ensure that the most appropriate architectures for science are developed while application and software researchers gain a better understanding of future systems to get a head start in developing software and models to take advantage of the new capabilities. Thus Research and Evaluation Prototype 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 2012, the Research and Evaluation Prototypes activity will support efforts to explore architectures on the path toward exascale computing and pursue multiple paths to overcoming key barriers. These efforts will be closely coordinated with FY 2012 investments in the NNSA to ensure complementary and efficient use of resources. These efforts will also be tightly coupled with research supported by the Applied Mathematics activity, the Computer Science activity, the SciDAC Co-Design Centers supported by the Computational Partnerships activity, and experience at the computing facilities.

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request

New Research and Evaluation Prototype investments will engage industry and leverage investments in the private sector as a partner in development. This is critical to ensuring that the demands of high performance computing for science and engineering are incorporated into commercial offerings and exascale capability. This effort will also provide the U.S. research community—from universities, national laboratories, and industry—with early access to emerging designs so as to inform their research efforts. Planned Research and Evaluation investments include partnerships with computer vendors to design the next generations of systems for science as well as focused partnerships to develop critical technologies:

- Initiation of two architectural development partnerships with U.S. computer vendors to develop hardware and low level software architectures which enable the creation of high performance scientific applications for these computers, as well as the smaller scale versions that will be ubiquitous in the scientific infrastructure; and
- Partnerships with U.S. memory vendors to deliver high bandwidth, power efficient memory technology for 2015 and 2018 computer systems. This is critical because the current commercial memory roadmaps indicate that memory power requirements will dominate the power budgets for computers targeted at scientific applications. Unless we invest in this area, the power requirements for exascale systems would be untenable. Low power, high speed memory with 10-100 fold improvement over current commercial offerings is required for science applications.

As the ASCAC report on exascale emphasized, "[f]or the information technology industry, leadership in hardware and software innovation for exascale computing will be accompanied by leadership in the full spectrum of computing from embedded computers and handheld mobile devices to laptop and desktop servers to departmental servers, all of which includes the use of high levels of intra-chip multithreaded parallelism. There is a critical need to find co-designed hardware and software solutions in this full spectrum (e.g., the power challenges for mobile devices bear a lot of similarity to those for nodes in an exascale systems) therefore whichever country leads in exascale computing will also enjoy technology leadership in the IT sector in general."[a]

High Performance Network Facilities and Testbeds	**29,982**	**34,500**

The costs for ESnet are dominated by operations which include refreshing hardware, such as switches and routers, on the schedule needed to ensure the 99.999% reliability that is required for large scale scientific data transmission.

The FY 2012 budget request supports operation of ESnet to continue to advance the next generation network capability that is critical to DOE applications and facilities. Building on the Recovery Act-supported Advanced Networking Initiative, ESnet will begin to deliver 100 gigabit per second (Gbps) connections to SC laboratories in FY 2012, with a goal of achieving 1,000 Gbps connectivity in 2016.

---

[a] *Opportunities and Challenges of Exascale Computing*, Advanced Scientific Computing Advisory Committee Subcommittee page 11, (http://www.science.doe.gov/ascr/ASCAC/Reports/Exascale-Subcommittee-Report.pdf).

	(dollars in thousands)	
	FY 2010 Current Appropriation	FY 2012 Request

	FY 2010	FY 2012
Achieved Operating Hours	8,760	N/A
Planned Operating Hours	8,760	8,760
Optimal Hours	8,760	8,760
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	0.01%	0.01%
Number of Users[a]	N/A	N/A

	FY 2010 Current Appropriation	FY 2012 Request
Small Business Innovative Research (SBIR)/Small Business Technology Transfer (STTR)	**0**	**7,464**

In FY 2010, $5,530,000 and $664,000 were transferred to the SBIR and STTR programs respectively. The FY 2012 amount shown is the estimated requirement for continuation of the SBIR and STTR programs.

	FY 2010 Current Appropriation	FY 2012 Request
Total, High Performance Computing and Network Facilities	**229,654**	**291,567**

Explanation of Funding Changes

	FY 2012 vs. FY 2010 Current Approp. ($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.	+2,900
Leadership Computing Facilities	
▪ Leadership Computing Facility at ORNL	
Increase supports site preparations and acquisition of a prototype hybrid multi-core computing system for research and tool development.	+7,212
▪ Leadership Computing Facility at ANL	
The increase supports upgrade of the facility to a 10 petaflop IBM Blue Gene/Q.	+20,000
Total, Leadership Computing Facilities	**+27,212**

[a] The ESnet is a high performance scientific network that connects DOE facilities to researches around the world and it is therefore not possible to estimate users.

	FY 2012 vs. FY 2010 Current Approp. ($000)
Research and Evaluation Prototypes	
The increase will support significant new research and development partnerships to advance exascale relevant prototypes and critical technologies that meet the needs of science applications.	+19,819
High Performance Network Facilities and Testbeds	
The increase will enable ESnet to plan for the installation and operation of a dedicated optical network to meet the growing requirements for DOE applications and facilities.	+4,518
SBIR/STTR	
Increase in SBIR/STTR due to increase in operating expenses.	+7,464
Total Funding Change, High Performance Computing and Network Facilities	**+61,913**

Supporting Information

Operating Expenses, Capital Equipment, and Construction Summary

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
Operating Expenses	373,899	440,600
Capital Equipment	9,300	25,000
Total, Advanced Scientific Computing Research	383,199	465,600

Funding Summary

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
Research	169,529	217,300
Scientific User Facility Operations	213,670	248,300
Total, Advanced Scientific Computing Research	383,199	465,600

Scientific User Facility Operations

(dollars in thousands)

	FY 2010 Current Appropriation	FY 2012 Request
NERSC	54,900	57,800
OLCF	86,788	94,000
ALCF	42,000	62,000
ESnet	29,982	34,500
Total, Scientific User Facility Operations	213,670	248,300

Facilities Users and Hours

	FY 2010 Current Appropriation	FY 2012 Request
NERSC		
Achieved Operating Hours	8,585	N/A
Planned Operating Hours	8,585	8,585
Optimal Hours	8,585	8,585
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	3,100	3,500

	FY 2010 Current Appropriation	FY 2012 Request
ESnet		
Achieved Operating Hours	8,760	N/A
Planned Operating Hours	8,760	8,760
Optimal Hours	8,760	8,760
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	0.01%	0.01%
Number of Users [a]	N/A	N/A
OLCF		
Achieved Operating Hours	7,008	N/A
Planned Operating Hours	7,008	7,008
Optimal Hours	7,008	7,008
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	550	625
ALCF		
Achieved Operating Hours	7,008	N/A
Planned Operating Hours	7,008	7,008
Optimal Hours	7,008	7,008
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	200	350
Total		
Achieved Operating Hours	31,361	N/A
Planned Operating Hours	31,361	31,361
Optimal Hours	31,361	31,361
Percent of Optimal Hours	100%	100%
Unscheduled Downtime	1%	1%
Number of Users	3,850	4,475

[a] The ESnet is a high performance scientific network that connects DOE facilities to researches around the world and it is therefore not possible to estimate users.

Scientific Employment

	FY 2010 Current Appropriation	FY 2012 Request
# University Grants	210	215
Average Size	224,000	247,000
# Laboratory Projects	180	181
# Graduate Students (FTEs)	488	527
# Permanent Ph.D.s (FTEs)	756	807
# Other (FTEs)	246	288