## Advanced Scientific Computing Research

### Overview

The Advanced Scientific Computing Research (ASCR) program's mission is to advance applied mathematics and computer science; deliver, in partnership with disciplinary science, the most advanced computational scientific applications; advance computing and networking capabilities; and develop, in partnership with the research community, including U.S. industry, future generations of computing hardware and tools for science. The strategy to accomplish this has two thrusts: developing and maintaining world-class computing and network facilities for science; and research in applied mathematics, computer science and advanced networking.

ASCR has a strong track record in scientific computing and innovation. For example, researchers at the Department of Energy (DOE) Leadership Computing Facilities (LCFs) have achieved breakthrough scientific and technological accomplishments that would not have been feasible without petascale scientific computing capabilities. The LCFs have enabled discoveries in functional materials, fundamental studies of turbulence in chemically reacting systems, improved understanding of climate change, and advances in the understanding of the physical properties of matter, such as quark-gluon nature of nuclear matter. Researchers using ASCR facilities have modeled 3-D full-core reactor neutron transport to predict the behavior of novel nuclear fuels; conducted 3D turbulent combustion simulations of hydrocarbons to increase fuel efficiency; made U.S. airplane engines quieter, more fuel efficient and less polluting; made long haul trucks more energy efficient in record time; simulated ice formation in million-molecule water droplets to reduce the wind turbine downtime in cold climates; and are developing novel materials for extreme energy environments.

The Office of Science, through ASCR, and the National Nuclear Security Administration (NNSA) have partnered to make strategic investments in hardware, methods, and critical technologies to maintain U.S. competitiveness in high performance computing (HPC). These efforts are linked with investments to advance data-intensive science and to effectively use the massive scientific data generated by DOE's unparalleled suite of scientific user facilities and large-scale collaborations. By investing in both next-generation computing and data-intensive science, the ASCR program will enable the community of users to improve and shorten industrial design processes; design advanced materials; better understand dark matter and dark energy; explore possibilities for dramatically increasing fuel efficiency while lowering emissions; design advanced nuclear reactors that are modular, safe, and affordable; improve accuracy of climate predictions; predict and investigate how to control the behavior of fusion plasmas; and calculate the subatomic interactions that determine nuclear structure.

### Highlights of the FY 2015 Budget Request

Capable exascale computing, with a hundred to thousand fold improvement in true application performance over today's systems, is the next frontier of development in High Performance Computing (HPC), extending capability significantly beyond today's petascale computers to address the next generation of scientific, engineering, and large-data problems. The goal of the exascale computing effort in SC is to provide the forefront computing resources needed to meet and advance the Department's science missions into the foreseeable future. This will require major new technology advances, the most important of which involve advances in parallelism and energy efficiency that are needed for scalable computing systems capable of sustained exaflop performance with acceptable power requirements. Research investments will impact computing at all scales from the largest scientific computers and data farms to department-scale computing to home computers and laptops.

The investment strategy has three components:

- Conduct research, development, and design efforts in hardware, software, and mathematical technologies leading toward capable exascale systems.
- Prepare today's scientific and data-intensive computing applications to migrate to and take full advantage of the emerging technologies from the research, development, and design efforts and eventually, exascale systems.

- Acquire and operate increasingly capable computing systems, starting with multi-petaflop machines that incorporate emerging technologies from research investments.

Mathematical, Computational, and Computer Sciences Research

Many of SC's user facilities and large experiments, such as light and neutron sources and the Large Hadron Collider (LHC), are moving to a distributed work environment where real-time interactions with instruments and simulations help to efficiently use facilities that often are located remotely from the control center.[a] Furthermore, experiments and simulations are often deeply intertwined as simulations help design large-scale experiments and data from experiments are used to validate or initialize simulations. The volume and complexity of data generated have increased such that a focused effort is required to develop theories, tools, and technologies to manage data from generation through integration, transformation, analysis, and visualization; to capture the historic record of the data; and to archive and share it. This request supports ASCR efforts in data-intensive science, and provides funding for a community of applied mathematicians and computer scientists to address end-to-end data management challenges.

This activity also develops hardware and software technologies and scalable system designs to broaden the impact of data-intensive science. To address the capable exascale challenge, this activity will support application development, co-design centers; research and development of software, tools, and middleware; and applied mathematics methods that address the challenges of both exascale and data-intensive science.

Software, tools, and methods from both of these efforts will be used by the Scientific Discovery through Advanced Computing (SciDAC) partnerships to more efficiently use the current and immediate next generation high performance computing facilities.

High Performance Computing and Network Facilities

Achieving a combined capability of several hundred petaflops at the LCFs requires technological advances in both hardware and software. In addition, Research and Evaluation Prototypes (REP) will support LCF-specific non-recurring engineering efforts to incorporate custom features that meet the Department's mission requirements. REP will also expand efforts in exascale component technology research and development, system engineering and integration leading to the design and development of future high performance computing (HPC) systems including prototype test beds for demonstrating the feasibility of building future full exascale systems and the exascale systems themselves.

The National Energy Research Scientific Computing Center (NERSC) will acquire the NERSC-8 supercomputer in FY 2015, which will expand the capacity of the facility by 10-40 petaflops to address emerging scientific needs. Also in FY 2015, NERSC will relocate to the Computational Research and Theory (CRT) building at Lawrence Berkeley National Laboratory.

Experienced computational scientists who assist a wide range of users in taking effective advantage of the advanced computing resources are critical assets at both the LCFs and NERSC. To address this DOE mission need, the LCFs and NERSC will work collaboratively to develop a post-doctoral training program for high end computational science and engineering.

As the Energy Science Network (ESnet) gains production experience with the 100 gigabit per second (Gbps) optical network, the need for next generation optical networking equipment has become clear. DOE is coordinating with other federal agencies to ensure the availability of next generation optical networking equipment from domestic sources. The outcomes of these efforts will help ESnet to keep pace with the continuing growth of scientific traffic from DOE's scientific user facilities.

[a] http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/ASCAC_Data_Intensive_Computing_report_final.pdf, http://science.energy.gov/~/media/ascr/pdf/research/scidac/ASCR_BES_Data_Report.pdf, and http://science.energy.gov/~/media/ascr/pdf/research/scidac/ASCR_BES_Data_Report.pdf

**Advanced Scientific Computing Research**
**Funding ($K)**

| | FY 2013 Current | FY 2014 Enacted | FY 2014 Current | FY 2015 Request | FY 2015 vs. FY 2014 Enacted |
|---|---|---|---|---|---|
| **Mathematical, Computational, and Computer Sciences Research** | | | | | |
| Applied Mathematics | 43,341 | 49,500 | 49,500 | 52,155 | +2,655 |
| Computer Science | 44,299 | 54,580 | 54,580 | 58,267 | +3,687 |
| Computational Partnerships | 41,971 | 46,918 | 46,918 | 46,918 | 0 |
| Next Generation Networking for Science | 11,779 | 15,931 | 15,931 | 19,500 | +3,569 |
| SBIR/STTR | 0 | 5,518 | 5,518 | 6,035 | +517 |
| **Total, Mathematical, Computational, and Computer Sciences Research** | **141,390** | **172,447** | **172,447** | **182,875** | **+10,428** |
| | | | | | |
| **High Performance Computing and Network Facilities** | | | | | |
| High Performance Production Computing | 62,000 | 65,605 | 65,605 | 69,000 | +3,395 |
| Leadership Computing Facilities | 146,000 | 160,000 | 160,000 | 184,637 | +24,637 |
| Research and Evaluation Prototypes | 24,000 | 37,784 | 37,784 | 57,934 | +20,150 |
| High Performance Network Facilities and Testbeds | 31,610 | 32,608 | 32,608 | 35,000 | +2,392 |
| SBIR/STTR | 0 | 9,649 | 9,649 | 11,554 | +1,905 |
| **Total, High Performance Computing and Network Facilities** | **263,610** | **305,646** | **305,646** | **358,125** | **+52,479** |
| **Total, Advanced Scientific Computing Research** | **405,000** | **478,093** | **478,093** | **541,000** | **+62,907** |

SBIR/STTR funding:

- FY 2013 Transferred: SBIR $11,312,000 and STTR $1,466,000
- FY 2014 Projected: SBIR $13,272,000 and STTR $1,895,000
- FY 2015 Request: SBIR $15,457,000 and STTR $2,132,000

**Advanced Scientific Computing Research**
**Explanation of Major Changes ($K)**

| | FY 2015 vs. FY 2014 Enacted |
|---|---|
| **Mathematical, Computational, and Computer Sciences Research:** Research will focus on the linked challenges of capable exascale and data-intensive science. There are increases for core research efforts in Applied Mathematics, Computer Science, and Next Generation Networking for Science for this purpose. The goal of these efforts is to develop new and improved, cross-disciplinary tools to manage and analyze massive scientific data, to initiate research and development of next generation storage technologies, and to increase support for research in managing on chip parallelism, fault tolerance and algorithm resilience, software approaches to energy management, and workflow and software environment tools. | +10,428 |
| **High Performance Computing and Network Facilities:** Increase supports the NERSC move and upgrade; lease costs, increased power costs, and preparations at the LCFs to support 75-200 petaflop upgrades at each facility; and ESnet will coordinate with efforts in other agencies to develop next generation optical networking equipment. Research and Evaluation Prototypes will expand investments in critical technologies and system integration for exascale and will support LCF specific research, design, and engineering efforts to ensure user facility upgrades meet the Department's mission requirements. | +52,479 |
| **Total, Advanced Scientific Computing Research** | **+62,907** |

**Basic and Applied R&D Coordination**

Coordination across disciplines and programs is a cornerstone of the ASCR program. Partnerships within SC are mature and continue to advance the use of high performance computing and scientific networks for science. A growing area of collaboration will be in the area of data-intensive science. ASCR continues to have a strong partnership with NNSA that is essential to achieving the Department's goals for exascale computing. In April 2011, ASCR and NNSA strengthened this partnership by signing a memorandum of understanding for collaboration and coordination of exascale research within the Department. Areas of mutual interest between ASCR and the DOE technology programs, particularly the Office of Electricity Delivery and Energy Reliability and the Office of Nuclear Energy, are applied mathematics for the optimization of complex systems, control theory, and risk assessment. Through the National Information Technology Research and Development Subcommittee of the National Science and Technology Council's Committee on Technology, the interagency networking and information technology R&D coordination effort, ASCR also coordinates with programs across the Federal Government.

**Program Accomplishments**

*Two ASCR Data Technologies Win R&D 100 Awards.* Among the winners of the 2013 R&D 100 Awards are ADIOS and OSCARS. ADIOS is the Adaptable I/O System for Big Data, developed by a team from Oak Ridge National Laboratory (ORNL), Georgia Tech, and Rutgers University and supported by ASCR through both Computer Science and SciDAC working closely with the LCF staff. ADIOS is an I/O middleware package that has shown great promise helping codes scale up with improved performance. OSCARS, the On-demand Secure Circuits and Reservation System, is a software service that creates dedicated bandwidth channels to move massive, time-critical data sets around the world, the development of which was led by the Next Generation of Networking for Science research program and implemented by Department of Energy's ESnet (Energy Sciences Network).

*Understanding the Long-Term Behavior of Geologically Sequestered Carbon Dioxide.* Computer models running at NERSC are generating data that nearly match the quality of images taken from experiments at the SC/Basic Energy Sciences Energy Frontier Research Center for Nanoscale Control of Geologic Carbon. These data are used to predict the fate of $CO_2$ trapped in the subsurface over long periods of time. The key to the success of the modeling effort – and the reason massive computation is required – is its ability to simultaneously represent both movements and chemical reactions that take place at microscopic levels as the $CO_2$ migrates through porous rock structures. The ability to capture subsurface behavior in a computer model is also relevant to several other key DOE missions, such as hydrocarbon recovery, legacy waste stewardship, and high-level waste isolation.

*Increasing Fuel Efficiency in the Near Term*. Through Oak Ridge's High Performance Computing (HPC) Industrial Program, engineers developed the first 3-D Computational Fluid Dynamics optimization of the complete underhood (UH3D) package of existing cars for an American automobile manufacturer. With UH3D, researchers were able to conduct engine bay analysis for the first time with the required number of design variables and operating conditions for a true design optimization. Additionally the new process will deliver more robust cooling system designs earlier in the new car development process reducing the number of costly physical prototypes. The results provided important return-on-investments justification for a significant upgrade to the company's in-house computing resources for similar projects.

*Optimizing the Grid in Real-time*. The production, distribution, storage, and use of electrical energy are undergoing significant changes. Demand and production patterns are being altered radically by the advent of "smart grids," renewable generation, hybrid electric vehicles, and storage technologies. The mathematical descriptions of such systems raise numerous challenges, ranging from multiple spatial and temporal variables to uncertainty in future operating conditions. For example, the Illinois grid contains approximately 2,000 transmission nodes, 2,500 transmission lines, 900 demand nodes, and 300 generation nodes; this translates into billions of variables and constraints once the uncertainty in the supply is taken into account. New algorithms have been developed that efficiently use up to 96% of the processors on the Argonne LCF and enabled simulation of wind energy in the Illinois power grid and energy market in "real-time" (within hours). This work demonstrated that certain classes of power grid problems, namely, energy dispatch problems, can significantly benefit from existing and emerging high-performance computing architectures. The work also addressed the problem of incorporating renewable energy sources into the U.S. power grid without increasing reserve units or degrading grid performance. Their work showed that even 20% wind penetration can be accommodated without significant reserve increase.

*Taking wind energy to new heights.* The amount of global electricity supplied by wind, the world's fastest growing energy source, is expected to increase from 2.5% to as much as 12% by 2020. Under ideal conditions, wind farms can be three times more efficient than coal plants, but going wherever the wind blows is not always so easy. Rather than risk investing in wind turbines that might freeze during icy, cold seasons at high latitudes or at high altitudes, these regions simply don't adopt wind power—despite powerful gusts of renewable energy nipping at their fingertips. Many turbines already located in colder regions rely on small heaters in the blades to melt ice, but these heaters drain 3–5% of the energy the turbine is producing. To better understand the underlying physical changes during ice formation on various surfaces researchers used nearly 40 million hours on Titan at Oak Ridge Leadership Computing Facility (OLCF) to simulate hundreds of millions of water molecules freezing in slow motion. These simulations replicated experimental results and deepened the understanding of freezing at the molecular level. Using visualizations developed by OLCF staff, the researchers studied the nucleation (or initial budding) of ice molecules among the millions of water molecules. These detailed simulations will help experimental researchers reduce the number of time-consuming and costly physical experiments in their goal to open cold climates to renewable power.

**Advanced Scientific Computing Research**
**Mathematical, Computational, and Computer Sciences Research**

**Description**

The Mathematical, Computational, and Computer Sciences Research subprogram supports research activities to effectively use the current and future generations of DOE's computer 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 delivers:

- new mathematics required to more accurately model systems involving processes taking place across a wide range of time and length scales;
- software, tools, and middleware to efficiently and effectively harness the potential of today's high performance computing systems and advanced networks for science and engineering applications;
- operating systems, data management, analyses, representation model development, user interfaces, and other tools required to make effective use of future-generation supercomputers and the data sets from current and future scientific user facilities;
- computer science and algorithm innovations that increase the energy efficiency of future-generation supercomputers;
- networking and collaboration tools to make scientific resources readily available to scientists, in university, national laboratory, and industrial settings.

The research program will develop methods, software, and tools to use commercial products for HPC systems for data-intensive and computational science. This requires a focus on concurrency, data movement, energy management, and resiliency to address individual chip failures and to improve HPC usability up to the exascale and exabytes.

Deriving scientific insights from vast amounts of raw data will require a focused research effort that will develop the necessary theories, tools, and technologies to manage the full data lifecycle from generation or collection through integration, transformation, analysis, and visualization, to capturing the historic record of the data and archiving, and sharing it. One result of this research effort will be a community of applied mathematicians and computer scientists available to address the Department's end-to-end data management challenges.

<u>Applied Mathematics</u>

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.

<u>Computer Science</u>

The Computer Science activity supports research on extreme-scale computing and extreme-scale data. Industry reports indicate that because of power constraints, data movement, rather than computational operations, will be the constraining factor for future systems. Memory per core is expected to decline sharply due to power requirements and the performance growth of storage systems will continue to lag behind the computational capability of the systems. Multi-level storage architectures that span multiple types of hardware are anticipated and require the activity to support research that develops new approaches to run-time data management and analysis.

A fundamental challenge is developing science applications to take advantage of technology advances such as multicore chips and specialized accelerator processors. This will require developing system software (operating systems, file systems, compilers, and performance tools) with more dynamic behavior than historically developed to deal with time-varying power and resilience requirements. Substantial innovation is needed to provide essential system software functionality in a timeframe consistent with the anticipated availability of hardware.

Computational Partnerships

The Computational Partnerships activity supports the SciDAC program, which accelerates progress in scientific computing through partnerships among applied mathematicians, computer scientists, and scientists in other disciplines. 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. SciDAC applications include climate science, fusion research, high energy physics, nuclear physics, astrophysics, materials science, chemistry, and accelerator physics.

SciDAC focuses on the high end of high performance computational science and engineering and addresses two challenges: to broaden the community and thus the impact of high performance computing, particularly to address the Department's missions and to ensure that further progress at the forefront is enhanced rather than curtailed by advances in computing hardware, most pressingly, the emergence of hybrid, multi-core architectures. SciDAC has helped U.S. industry to use computing to improve competitiveness.

Next Generation Networking for Science

ASCR has played a leading role in the development of the high-bandwidth networks connecting researchers to facilities, data, and one another. ASCR-supported researchers helped establish critical protocols on which the internet is based. Next Generation Networking for Science research makes possible international collaborations such as the Large Hadron Collider and underpins virtual meeting and other commercial collaboration tools. These research efforts build upon results from Computer Science and Applied Mathematics to develop integrated software tools and advanced network services to use new capabilities in ESnet to advance DOE missions.

## Activities and Explanation of Change

| FY 2014 Enacted | FY 2015 Request | Explanation of Changes FY 2015 vs. FY 2014 Enacted |
|---|---|---|
| **Applied Mathematics** | | |
| Significant innovation in applied mathematics is needed to realize the potential of next generation high performance computing systems. The Applied Mathematics portfolio will shift toward investments aimed at addressing these critical research challenges, as well as associated challenges in complex systems and data–intensive science. Energy management, data movement, and resiliency research will be emphasized. | High-fidelity modeling and simulation will require a number of new algorithmic techniques and strategies supported by this activity, including: advanced solvers for large linear and nonlinear systems, time integration schemes, multi-physics coupling, methods that use asynchrony or randomness, adaptively, algorithmic resilience, and strategies for reducing global communications. | Increased research efforts support the Department's efforts in capable exascale and data-intensive science. These efforts will develop scalable mathematical and statistical models, algorithms, and methods for the representation, analysis, and understanding of extreme-scale data from scientific simulations and experiments. |
| **Computer Science** | | |
| Significant innovation in computer science is needed to realize the potential of next generation high performance computing (HPC) systems and other scientific user facilities. There will be continued emphasis on data-intensive science challenges with particular attention to the intersection with exascale computing challenges and the unique needs of DOE scientific user facilities including data management. There will also be expanded efforts in tools, user interfaces, the high performance computing software stack, and visualization and analytics. These efforts are essential to ensure DOE mission applications are able to use commercially available HPC hardware. | To achieve the full potential of exascale computing, a software stack must be developed that includes new programming models and metrics for evaluating system status. This activity will support software efforts that span the spectrum from low-level, operational software to high-level, application development environments. More specifically, it will include operating systems, runtimes for scheduling, memory management, file systems, and performance monitoring. Also included are power management and resilience strategies, computational libraries, compilers, programming models, and application frameworks. Scalability, programmability, resilience, and code portability will be emphasized to promote ease of use, reliability, accommodation of legacy code, and pathways to future development beyond exascale. | Increased research efforts support the Department's efforts in capable exascale and data-intensive science. These efforts will focus on in situ data management, analysis and visualization, new I/O subsystems, and new multi-level storage system software. |

| FY 2014 Enacted | FY 2015 Request | Explanation of Changes FY 2015 vs. FY 2014 Enacted |
|---|---|---|
| **Computational Partnerships** | | |
| The SciDAC institutes continue to play a key role in assisting DOE mission critical applications to effectively use the ASCR production and leadership computing facilities. The strategic partnerships with the other Office of Science programs will continue to address their specific needs as they move toward larger data sets and more complex computing systems.<br><br>This activity will also focus on the current set of co-design centers that partner DOE mission applications with forefront researchers and computing vendors. These efforts will inform core research efforts in applied mathematics and computer science as well as the computing resources for the next generation of scientific user facilities. | The Scalable Data Management Analysis and Visualization (SDAV) Institute and data-intensive co-design center will continue to support efforts in the Department's efforts in data-intensive science. The role of this activity is to develop robust tools and software to manage and analyze massive data with SDAV focused on the near term and co-design focused on emerging hardware. | |
| **Next Generation Networking for Science** | | |
| With the production deployment of 100 gigabit per second (Gbps) technologies, research will continue to focus on developing networking software, middleware, and hardware that delivers 99.999% reliability while allowing the successful products of prior research to transition into operation. These investments are increasingly important as ESnet expands production use of very high-throughput and optical technologies. | Research will focus on the challenges of moving, sharing, and validating massive quantities of data from DOE scientific user facilities and large scale collaborations via high speed optical networks. This includes the challenges in building, operating, and maintaining the network infrastructure over which these data pass. | Increased research efforts support the Department's efforts in exascale and data-intensive science. This activity will focus on integrating the SC facilities with computing resources and mechanisms for discovering resources and data in a globally distributed computing environment. |
| **SBIR/STTR** | | |
| SBIR/STTR funding is set at 3.2% of non-capital funding in FY 2014. | In FY 2015, SBIR/STTR funding is set at 3.3% of non-capital funding. | The SBIR/STTR amount is adjusted to mandated percentage for non-capital funding. |

## Advanced Scientific Computing Research
## High Performance Computing and Network Facilities

### Description

The High Performance Computing and Network Facilities subprogram delivers forefront computational and networking capabilities. These include high performance production computing at the National Energy Research Scientific Computing Center (NERSC) at LBNL and Leadership Computing Facilities (LCFs) at ORNL and ANL. 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 Research and Evaluation Prototypes activity invests in research and development that will play a critical role in delivering world-leading capabilities and achieving the Department's exascale computing goals.

Allocations on ASCR facilities provide critical resources for the scientific community, including industry and other agencies, following the peer reviewed, public access model used by other SC scientific user facilities. ASCR facilities provide a testbed for U.S. industry to scale and then validate code performance to optimize in-house HPC investments.

The Research and Evaluation Prototypes activity addresses the challenges of next generation computing systems. By actively partnering with the research community, including industry, on the development of technologies that enables next-generation machines, ASCR ensures that commercially available architectures serve the needs of the scientific community. Coupling this activity to the co-design centers ensures that application and software researchers can gain a better understanding of future systems to get a head start in developing software and models to take advantage of the new capabilities. Research and Evaluation Prototypes prepares researchers to effectively use the next generation of scientific computers and seeks to reduce risk for future major procurements.

<u>High Performance Production Computing</u>

This activity supports NERSC, which delivers high-end production computing services for the SC research community. Annually, approximately 5,000 computational scientists in about 500 projects use NERSC to perform scientific research across a wide range of disciplines including astrophysics, chemistry, climate modeling, materials, high energy and nuclear physics, and biology. NERSC users come from nearly every state in the U.S., with about 65% based in universities, 25% in DOE laboratories, and 10% in other government laboratories and industry. NERSC's large and diverse user base requires an agile support staff to aid users entering the high performance computing arena for the first time, as well as those preparing codes to run on the largest machines available at NERSC and other SC computing facilities.

NERSC is a vital resource for the SC research community and it is consistently oversubscribed, with requests exceeding capacity by a factor of 3–10. This gap between demand and capacity exists despite upgrades to the primary computing systems approximately every 3 years. NERSC regularly gathers requirements from SC programs through a robust process that informs NERSC upgrade plans. These requirements activities are also vital to planning for SciDAC and other ASCR efforts to prioritize research directions and inform the community of new computing trends, especially as the computing industry moves toward heterogeneous, multi-core computing.

<u>Leadership Computing Facilities</u>

The LCFs enable open scientific applications, including industry applications, to harness the potential of leadership computing to advance science and engineering. The era of petaflop science opened significant opportunities to dramatically advance research as simulations more realistically capture complex behavior in natural and engineered systems. The success of this effort is built on the gains made in Research and Evaluation Prototypes and ASCR research efforts. LCF staff operates and maintains forefront computing resources. One LCF strength is the staff support provided to Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects, ASCR Leadership Computing Challenge projects, scaling tests, early science applications, and tool and library developers. Support staff experience is critical to the success of industry partnerships to address the challenges of next-generation computing.

The Oak Ridge Leadership Computing Facility (OLCF) 27 petaflop system is one of the most powerful computers in the world for scientific research, according to the June 2013 Top 500 list. Through INCITE allocations, several applications, including combustion studies in diesel jet flame stabilization, simulations of neutron transport in fast-fission reactor cores, and groundwater flow in porous media, are running at the multi-petaflop scale. OLCF staff is sharing its expertise with industry to broaden the benefits for the Nation. For example, OLCF worked with industry to significantly reduce the need for costly physical prototypes and physical tests to design energy efficient cooling systems leading to increased fuel efficiency in vehicles.

The Argonne Leadership Computing Facility (ALCF) operates a 10 petaflop IBM Blue Gene Q (Mira) machine with relatively low electrical power requirements. The IBM Blue Gene/Q was developed through a joint research project with support from the NNSA, industry, 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. ALCF supports many applications, including molecular dynamics and materials, for which it is better suited than OLCF or NERSC. Through INCITE, ALCF also transfers its expertise to industry, for example, helping engineers understand the fundamental physics of turbulent mixing to transform product design and to achieve improved performance, life and efficiency.

The demand for 2014 INCITE allocations at the LCFs outpaced the available resources by a factor of three.

<u>Research and Evaluation Prototypes</u>

The next generation of computing hardware will present new challenges for science and engineering applications—most notably higher levels of concurrency involving billions of processing elements, effectively managing chip failures and silent errors, and power demands that will restrict memory usage. This activity supports research and development partnerships with vendors to influence and accelerate critical technologies for next-generation systems, system integration research, and development and engineering efforts that are coupled to application development to ensure Department applications are ready to make effective use of commercial offerings.

In addition, this activity partners with the NNSA to support research investments in non-recurring engineering, for near-term technology customization for the ASCR facilities.

<u>High Performance Network Facilities and Testbeds</u>

The Energy Sciences Network (ESnet) provides the national network and networking infrastructure connecting DOE science facilities and SC laboratories with other institutions connected to peer academic or commercial networks. The costs for ESnet are dominated by operations, including maintaining the fiber optic backbone and refreshing switches and routers on the schedule needed to ensure the 99.999% reliability required for large-scale scientific data transmission. Additional funds are used to support the growth in science data traffic and for testing and evaluation of new technologies and services that will be required to keep pace with the expected data volume. DOE is coordinating with other federal agencies to ensure the availability of next generation optical networking equipment from domestic sources. The outcomes of these efforts will help ESnet to keep pace with the continuing growth of scientific traffic from DOE's scientific user facilities.

## Advanced Scientific Computing Research
## High Performance Computing and Network Facilities

### Activities and Explanation of Changes

| FY 2014 Enacted | FY 2015 Request | Explanation of Changes FY 2015 vs. FY 2014 Enacted |
|---|---|---|
| **High Performance Production Computing** | | |
| Supports operation of the NERSC capability systems (NERSC-7) including power costs, lease payments, and user support. Also supports, as part of the NERSC-7 upgrade project, continued site preparations for the new NERSC facility on the LBNL campus. | Supports operation of the NERSC high-end capability systems (NERSC-7) including increased power costs, lease payments, and user support and a new post-doctoral training program for high-end computational science and engineering. | Increase supports operations, lease payments, and user support for NERSC including NERSC-7, which will more than double the capacity of NERSC but also increases the power requirement of NERSC. Also supports the move into the Computational Research and Theory building. Also supports initiation of a post-doctoral training program for high-end computational science and engineering. |
| **Leadership Computing Facilities** | | |
| Supports operation and allocation, through INCITE and ALCC, of the upgraded 20 petaflop OLCF and 10 petaflop ALCF. This includes lease payments, power, and user support. | Supports operation and allocation, through INCITE and ALCC, of the 27 petaflop Titan system at the OLCF and 10 petaflop Mira system at the ALCF. This includes lease payments, power, and user support. Also supports preparations at the LCFs to support 75-200 petaflop upgrades at each facility and a new post-doctoral training program for high-end computational science and engineering. | Increase also supports operations, lease payments, additional power costs, and user support for the 27 petaflop system at the OLCF and 10-petaflop machine at the ALCF and preparations for the next upgrades. Also supports initiation of a post-doctoral training program for high-end computational science and engineering. |
| Leadership Computing Facility at ANL: $67,000 | $80,320 | +$13,320 |
| Leadership Computing Facility at ORNL: $93,000 | $104,317 | +$11,317 |

| FY 2014 Enacted | FY 2015 Request | Explanation of Changes<br>FY 2015 vs. FY 2014 Enacted |
|---|---|---|
| **Research and Evaluation Prototypes** | | |
| Research and Evaluations Prototypes (REP) will continue joint investments between ASCR and NNSA's Advanced Simulation and Computing (ASC) program in the *Fast Forward* critical technology effort, whose primary objective is to form partnerships with industry to accelerate the R&D of critical component technologies needed for extreme-scale computing. The selected technologies have the potential to impact low-power embedded, cloud/datacenter, and mid-range HPC applications, thus ensuring that DOE/NNSA investment furthers a sustainable software/hardware ecosystem supported by applications across the HPC market and the broader IT industry. As part of the *Fast Forward* effort, REP will initiate investments exploring early exascale system design and analysis of interconnect design trade-offs. | As follow-ons to the previous research efforts in critical technologies and system interconnect, REP will competitively select teams to develop system designs suitable for next-generation platforms. In addition, it will fund the development of prototypes based on the results from the *Fast Forward* program's investments in critical crosscutting technology research in areas such as processors, memory subsystems, network interfaces and the interconnection network. | Increase supports the initiation of system design and prototype development efforts. Increase also supports LCF specific non-recurring engineering efforts to ensure the planned LCF upgrades meet the Department's mission requirements. |
| **High Performance Network Facilities and Testbeds** | | |
| ESnet will operate the network infrastructure to support critical DOE science applications and SC facilities. ESnet extends deployment of 100 Gbps production network by connecting remaining SC laboratories at 100 Gbps speeds. | ESnet will operate the network infrastructure to support critical DOE science applications, SC facilities and scientific collaborations around the world through 100 Gbps production network. | DOE is coordinating with other federal agencies to ensure the availability of next generation optical networking equipment from domestic sources. These efforts will ensure ESnet keeps pace with continuing growth of scientific traffic from DOE's scientific user facilities and scientific collaborations. Also supports lighting of additional 100 GBS fiber to support interim traffic growth. |
| **SBIR/STTR** | | |
| SBIR/STTR funding is set at 3.2% of non-capital funding in FY 2014. | In FY 2015, SBIR/STTR funding is set at 3.3% of non-capital funding. | The SBIR/STTR amount is adjusted to mandated percentage for non-capital funding. |

**Advanced Scientific Computing Research**
**Performance Measures**

In accordance with the GPRA Modernization Act of 2010, the Department sets targets for, and tracks progress toward, achieving performance goals for each program. For more information, refer to the Department's FY 2013 Annual Performance Report. The following table shows the targets for FY 2013 through 2015.

| | FY 2013 | FY 2014 | FY 2015 |
|---|---|---|---|
| Performance Goal (Measure) | **ASCR Facility Operations—Average achieved operation time of ASCR user facilities as a percentage of total scheduled annual operation time** | | |
| Target | ≥ 90% | ≥ 90% | ≥ 90% |
| Result | Met | TBD | TBD |
| Endpoint Target | Many of the research projects that are undertaken at the Office of Science's scientific user facilities take a great deal of time, money, and effort to prepare and regularly have a very short window of opportunity to run. If the facility is not operating as expected the experiment could be ruined or critically setback. In addition, taxpayers have invested millions or even hundreds of millions of dollars in these facilities. The greater the period of reliable operations, the greater the return on the taxpayers' investment. | | |
| Performance Goal (Measure) | **ASCR Research—Discovery of new applied mathematics and computer science tools and methods that enable DOE applications to deliver scientific and engineering insights with a significantly higher degree of fidelity and predictive power** | | |
| Target | Accept and put into service 10 petaflop upgrades at Argonne and Oak Ridge Leadership Computing Facilities | Initiate at least four new teams to conduct fundamental computer science research and at least three new applied mathematics research awards that address issues of fault tolerance or energy management for next-generation computing systems | Conduct an external peer review of the three original co-design centers to document progress, impact, and lessons learned. |
| Result | Met | TBD | TBD |
| Endpoint Target | Develop and deploy high-performance computing hardware and software systems through exascale platforms. | | |

## Advanced Scientific Computing Research
### Capital Summary ($K)

| | Total | Prior Years | FY 2013 Current | FY 2014 Enacted | FY 2014 Current | FY 2015 Request | FY 2015 vs. FY 2014 Enacted |
|---|---|---|---|---|---|---|---|
| **Capital operating expenses** | | | | | | | |
| Capital equipment | n/a | n/a | 0 | 4,100 | 4,100 | 8,000 | +3,900 |

### Funding Summary ($K)

| | FY 2013 Current | FY 2014 Enacted | FY 2015 Request | FY 2015 vs. FY 2014 Enacted |
|---|---|---|---|---|
| Research | 165,390 | 204,713 | 234,774 | +30,061 |
| Scientific user facility operations | 239,610 | 258,213 | 288,637 | +30,424 |
| Other | 0 | 15,167 | 17,589 | +2,422 |
| **Total, Advanced Scientific Computing Research** | **405,000** | **478,093** | **541,000** | **+62,907** |

### Scientific User Facility Operations ($K)

| | FY 2013 Current | FY 2014 Enacted | FY 2015 Request | FY 2015 vs. FY 2014 Enacted |
|---|---|---|---|---|
| **NERSC** | $62,000 | $65,605 | $69,000 | +$3,395 |
| Achieved operating hours | N/A | N/A | N/A | |
| Planned operating hours | 8,585 | 8,585 | 8,585 | 0 |
| Optimal hours | 8,585 | 8,585 | 8,585 | 0 |
| Percent of optimal hours | 100% | 100% | 100% | |
| Unscheduled downtime percentage | 1% | 1% | 1% | |
| Number of users | 5,000 | 5,500 | 5,500 | 0 |

| | FY 2013 Current | FY 2014 Enacted | FY 2015 Request | FY 2015 vs. FY 2014 Enacted |
|---|---|---|---|---|
| **OLCF** | $85,000 | $93,000 | $104,317 | +$11,317 |
| Achieved operating hours | N/A | N/A | N/A | |
| Planned operating hours | 7,008 | 7,008 | 7,008 | 0 |
| Optimal hours | 7,008 | 7,008 | 7,008 | 0 |
| Percent of optimal hours | 100% | 100% | 100% | |
| Unscheduled downtime percentage | 2% | 1% | 1% | |
| Number of users | 1,200 | 1,300 | 1,300 | 0 |
| | | | | |
| **ALCF** | $61,000 | $67,000 | $80,320 | +$13,320 |
| Achieved operating hours | N/A | N/A | N/A | |
| Planned operating hours | 7,008 | 7,008 | 7,008 | 0 |
| Optimal hours | 7,008 | 7,008 | 7,008 | 0 |
| Percent of optimal hours | 100% | 100% | 100% | |
| Unscheduled downtime percentage | 2% | 1% | 1% | |
| Number of users | 900 | 1,000 | 1,000 | 0 |
| | | | | |
| **ESnet** | $31,610 | $32,608 | $35,000 | +$2,392 |
| Achieved operating hours | N/A | N/A | N/A | |
| Planned operating hours | 8,760 | 8,760 | 8,760 | 0 |
| Optimal hours | 8,760 | 8,760 | 8,760 | 0 |
| Percent of optimal hours | 100% | 100% | 100% | |
| Unscheduled downtime percentage | 0.01% | 0.01% | 0.01% | |
| Number of users[a] | N/A | N/A | N/A | 0 |
| **Total, Scientific User Facility Operations** | $239,610 | $258,213 | $288,637 | +$30,424 |
| Achieved operating hours | N/A | N/A | N/A | |
| Planned operating hours | 31,361 | 31,361 | 31,361 | 0 |
| Optimal hours | 31,361 | 31,361 | 31,361 | 0 |
| Percent of optimal hours (funding weighted) | 100% | 100% | 100% | |
| Unscheduled downtime percentage | 1% | 1% | 1% | |
| Number of users | 7,100 | 7,800 | 7,800 | 0 |

[a] ESnet is a high performance scientific network connecting DOE facilities to researchers around the world; user statistics are not collected.

**Scientific Employment**

| | FY 2013 Estimate | FY 2014 Estimate | FY 2015 Estimate | FY 2015 vs FY 2014 |
|---|---|---|---|---|
| Number of graduate students (FTEs) | 510 | 400 | 440 | +40 |
| Number of permanent Ph.D.'s (FTEs) | 725 | 650 | 720 | +70 |
| Other scientific employment (FTEs) | 255 | 220 | 245 | +25 |