Advanced Scientific Computing Research

Overview

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

During the past six decades, U.S. computing capabilities have been maintained through sustained research and the development and deployment of new computing systems with rapidly increasing performance on applications of major significance to government, industry, and academia. The Department of Energy (DOE) and its predecessor organizations have played a key role in that process–advancing national security, science, and industrial competitiveness. To maximize the benefits of high performance computing (HPC) in the coming decades, DOE will sustain and enhance its scientific, technological, and economic leadership position in HPC research, development, and deployment as part of a coordinated Federal strategy guided by four principles:

- Deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
- Foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.
- Draw upon the strengths of and seek cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.
- Develop a comprehensive technical and scientific approach to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

On July 29, 2015, an executive order established the National Strategic Computing Initiative (NSCI) to ensure a coordinated Federal strategy in HPC research, development, and deployment. DOE, along with the Department of Defense and the National Science Foundation, have been selected to co-lead the NSCI. Specifically, the DOE Office of Science (SC) and the DOE National Nuclear Security Administration (NNSA) are responsible for the execution of a joint program focused on advanced simulation through a capable exascale computing program, with an emphasis on sustained performance on science and national security mission applications and increased convergence between exascale and large-data analytic computing.

DOE will meet its NSCI assignment through the Exascale Computing Initiative (ECI), which began in FY 2016. The ECI, which is a partnership between SC and NNSA, will accelerate the research and development (R&D) to overcome key exascale challenges in parallelism, energy efficiency, and reliability, leading to deployment of exascale systems in the mid-2020s. The acceleration or advancement is defined as a hundred-fold increase in sustained performance over today's computing capabilities, enabling applications to address next-generation science, engineering, and data problems. The plan for the ECI has been reviewed by the interagency community and the Advanced Scientific Computing Advisory Committee (ASCAC).

In addition to underpinning DOE's mission in science, capabilities developed in the ECI will also support DOE's applied energy technology developments. ECI is critical to advancing energy technologies. The DOE's Quadrennial Technology Review (QTR), released in September 2015, identifies R&D opportunities across the six energy technology areas and documents a cross-cutting need for modeling, simulation, and analytics. ECI is a critical tool in support of advancing energy technologies.

In FY 2017, the ASCR portion of the Office of Science component of the ECI is contained in the Office of Science Exascale Computing Project (SC-ECP), which includes only the activities required for the delivery of exascale computers.

The scope of the SC-ECP has four focus areas:

- *Hardware Technology:* The Hardware Technology focus area supports vendor-based R&D activities required to deploy at least two exascale systems with diverse architectural features. Within this focus area, a node design effort targets component technologies needed to build exascale nodes, including the required software, while a system design effort performs the engineering and R&D activities required to build a full exascale computer and the required systems software;
- *System Software Technology:* The System Software Technology focus area spans low-level operational software to programming environments for high-level applications software development, including the software infrastructure to support large data management and data science for the DOE at exascale;
- *Application Development:* The Application Development focus area includes: extreme parallelism, reliability and resiliency, deep hierarchies of hardware processors and memory, scaling to larger systems, and data-intensive science; and
- *Exascale Systems:* The Exascale Systems focus area supports advanced system engineering development by the vendors needed to produce capable exascale systems. System procurement activities will be executed in coordination with each DOE HPC facility's existing system acquisition timelines. It also includes acquisition and support of prototypes and testbeds for the application, software, and hardware testing activities.

The SC-ECP will be managed following the principles of DOE Order 413.3B, *Program and Project Management for the Acquisition of Capital Assets*, which has been used by the Office of Science for the planning, design, and construction of all of its major projects, including the Leadership Computing Facilities at Argonne and Oak Ridge National Laboratories (ORNL) and NERSC at Lawrence Berkeley National Laboratory (LBNL). Computer acquisitions use a tailored version of Order 413.3B. The first four years of SC-ECP will focus on research in software (new algorithms and methods to support application and system software development) and hardware (node and system design). During the last six years of the SC-ECP, activities will focus on delivering application software, the system software stack, and hardware technologies that will be deployed in the exascale systems.

Overall project management for SC-ECP will be conducted via a Project Office that has been established at ORNL, which has considerable expertise in developing computational science and engineering applications and in managing HPC facilities, both for the Department and for other federal agencies. ORNL also has experience in managing distributed, large-scale scientific research projects, such as the Spallation Neutron Source project—a collaboration among six DOE national laboratories with a TPC of $1.4 billion. The Project Office is now initiating coordination among partners and developing the bases for pending Critical Decisions for the SC-ECP.

Highlights of the FY 2017 Budget Request

The FY 2017 Budget Request for ASCR continues support for the basic and applied research activities that support the broad scientific objectives of the Office of Science. The ASCR budget also implements the DOE responsibilities defined in the Administration's National Strategic Computing Initiative and supports the Department's Agency Priority Goal on high performance computing, by accelerating the delivery of a capable exascale computing system. A capable exascale computing system integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs, including data-intensive science. In the FY 2017 Budget Request, most funding for the Working Capital Fund (WCF) is transferred to Science Program Direction to establish a consolidated source of funding for goods and services provided by the WCF. CyberOne is still funded through program dollars in the Office of Science Safeguards and Security program. In FY 2016 and prior years, WCF costs were shared by SC research programs and Program Direction.

Mathematical, Computational, and Computer Sciences Research

With the initiation of the SC-ECP, research efforts that are on the critical path for the ECI have been shifted from the Mathematical, Computational, and Computer Sciences Research subprogram to the Exascale Computing subprogram.

As noted in the NSCI, the era of silicon-based microchips advancing in accordance with Moore's Law (feature sizes reducing by a factor of two approximately every two years) is nearing an end due to limits imposed by fundamental physics. ASCR will invest $12 million across research and facilities to understand the impacts these technologies may have on our applications. Beginning in FY 2017, the computer science and computational partnerships activities will invest $7 million to initiate new research efforts on technologies “Beyond Moore’s Law,” responding to the NSCI and recommendations made by the Secretary of Energy Advisory Board, to understand the challenges that these dramatically different technologies pose to DOE mission applications and to identify the hardware, software and algorithms that will need to be developed for DOE mission applications to harness these developing technologies.

The NSCI also drew attention to the need to increase the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development. Activities in applied mathematics, computer science, and next generation networking for science that are not included in the ECI provide the foundation for increasing the capability of the national HPC ecosystem. In FY 2017, these activities will continue to develop the methods, software, and tools to ensure DOE applications can fully exploit the most advanced computing systems available and use HPC systems for data-intensive and computational science at the exascale and beyond.

Software, tools, and methods developed by these core research efforts are used by the Scientific Discovery through Advanced Computing (SciDAC) computational partnerships. This allows the other scientific programs in the Office of Science to more effectively use the current and immediate next generation high performance computing facilities. In FY 2017, the SciDAC activity will be re-competed and expanded. The focus of the new SciDAC portfolio will be on developing the mission critical applications of the other Office of Science programs. These efforts will be informed by the research results emerging from the exascale computing initiative and will, whenever possible, incorporate the software, methods, and tools developed by that initiative. The increase in Computational Partnerships also initiates: a partnership with BES and BER to develop new tools and technologies for the BRAIN initiative, a partnership with NNSA on seismic simulation, and partnerships supporting the Administration’s Clean Energy Initiatives.

High Performance Computing and Network Facilities

With the initiation of the SC-ECP, research efforts that are on the critical path for the ECI have been shifted from the High Performance Computing and Network Facilities subprogram to the Exascale Computing Subprogram.

In FY 2017, the Leadership Computing Facilities (LCFs) will continue preparations for planned 75-200 petaflops (pf) upgrades at each site to be completed in the 2018-2019 timeframe. The Argonne LCF (ALCF) will also deploy an interim system in FY 2017 to transition ALCF users to the new many-core architecture being introduced by computer vendors in that time frame. Because these upgrades represent technological advances in both hardware and software, funds are included in Research and Evaluation Prototypes (REP) to continue supporting non-recurring engineering efforts for the ASCR facilities that incorporate custom features to meet the Department’s mission requirements.
The National Energy Research Scientific Computing Center (NERSC) will begin operation of the NERSC-8 supercomputer, named Cori, which will expand the capacity of the facility to approximately 30pf to address the continued increase in demand from Office of Science researchers. To keep pace with the growing demand for capacity computing to meet mission needs, the Department has begun planning for deployment of NERSC-9 in 2020, which will have three to five times the capacity of NERSC-8.

Given the significant external competition for trained workforce across the ASCR portfolio and the need to develop the workforce to support the objectives of the NSCI, the REP activity will continue to support the Computational Sciences Graduate Fellowship at $10,000,000 in FY 2017. 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, support also continues for a post-doctoral training program for high end computational science and engineering. In addition, the three ASCR HPC user facilities will continue coordinating efforts to quantify scientist’s computational requirements and to prepare their users for future architectures.

To support the new research effort within the computer science and computational partnerships activities, REP will support a small-scale testbed for technologies that are “Beyond Moore’s Law.” Given the increasing threat from cyber-attacks on federal resources and the expertise within the ASCR research community, the research and evaluation prototypes activity will also initiate a modest research effort in cybersecurity in FY 2017 with an emphasis on the unique challenges of the Department’s HPC facilities, which are not currently addressed by ongoing cyber-security R&D.

In FY 2017, the Energy Science Network (ESnet) will provide increases in bandwidth to address the growing data requirements of SC facilities, such as the Department’s light sources, neutron sources, and particle accelerators at CERN.

This includes upgrading high-traffic links to 400 gigabits per second (gbps). ESnet will also continue to extend science engagement efforts to solve the end-to-end network issues between DOE facilities and universities.

Exascale Computing

With the initiation of the SC-ECP, activities that are on the critical path for the ECI have been shifted to this new subprogram.

The primary goal of the ECI in SC is to provide the forefront computing resources needed to meet and advance the Department's science missions into the next decade. This will require major advances in technology, the most important of which are increased parallelism, energy efficiency, and reliability, which are needed for scalable use of these computing systems. Because DOE partners with HPC vendors to accelerate the development of commodity parts, these research investments will impact computing at all scales, ranging from the largest scientific computers and data centers to department-scale computing to home computers and laptops. The FY 2017 Request includes $154,000,000 for the SC-ECP.

The investment strategy for the ECI has five 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 exploit fully the capabilities of exascale systems by coordinating their development with the emerging technologies from the research, development, and design efforts.
- Partner with HPC vendors to accelerate the pace of implementation of technologies required for capable exascale computing.
- Acquire and operate increasingly capable computing systems, starting with hundred-plus petaflop machines that incorporate emerging technologies from research investments.
- Collaborate with other Federal agencies to ensure broad applicability and use of capable exascale computing across the US Government.

Within the FY 2017 Budget Request, ASCR supports the Department's ECI goal to significantly accelerate the development of capable exascale computing systems to meet national needs through the SC-ECP. Exascale computing systems, capable of at least one billion billion (1×10^{18}) calculations per second, are needed to advance science objectives in the physical sciences, such as materials and chemical sciences, high-energy and nuclear physics, climate and energy modeling, genomics and systems biology, as well as to support national security objectives and applied-energy research advances in DOE. Exascale systems' computational power is also needed for increasing data-analytic and data-intense applications across the set of DOE science programs and other Federal organizations that rely on large-scale simulations, e.g., the National Oceanographic and Atmospheric Administration and the National Institutes of Health. The importance of exascale computing to the DOE science programs is documented in previous and ongoing individual requirements reviews for each SC program office. Exascale computing is a central component of long-term collaboration between the SC's Advanced Scientific Computing Research (ASCR) program and the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing Campaign (ASC) program.

FY 2017 Crosscuts ($K)

	ECI
Advanced Scientific Computing Research	154,000

Advanced Scientific Computing Research
Funding ($K)

	FY 2015 Enacted	FY 2015 Current[a]	FY 2016 Enacted	FY 2017 Request[b]	FY 2017 vs FY 2016
Mathematical, Computational, and Computer Sciences Research					
Applied Mathematics	49,155	49,454	49,229	39,229	-10,000
Computer Science	55,767	55,259	56,848	39,296	-17,552
Computational Partnerships	46,918	43,996	47,918	45,596	-2,322
Next Generation Networking for Science	19,000	19,011	19,000	19,000	0
SBIR/STTR	5,830	0	6,181	7,733	+1,552
Total, Mathematical, Computational, and Computer Sciences Research	**176,670**	**167,720**	**179,176**	**150,854**	**-28,322**
High Performance Computing and Network Facilities					
High Performance Production Computing	75,605	75,905	86,000	92,145	+6,145
Leadership Computing Facilities	184,637	190,698	181,317	187,000	+5,683
Research and Evaluation Prototypes	57,329	53,298	121,471	17,890	-103,581
High Performance Network Facilities and Testbeds	35,000	35,790	38,000	45,000	+7,000
SBIR/STTR	11,759	0	15,036	16,291	+1,255
Total, High Performance Computing and Network Facilities	**364,330**	**355,691**	**441,824**	**358,326**	**-83,498**
Exascale Computing					
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	0	0	0	154,000	+154,000
Total, Advanced Scientific Computing Research	**541,000**	**523,411**	**621,000**	**663,180**	**+42,180**

SBIR/STTR funding:

- FY 2015 Enacted: SBIR $15,457,000 and STTR $2,132,000
- FY 2016 Request: SBIR $18,450,000 and STTR $2,767,000
- FY 2017 Request: SBIR $21,062,000 and STTR $2,962,000

---

[a] Reflects the transfer of Small Business Innovation/Technology Transfer Research (SBIR/STTR) funds within the Office of Science.

[b] A transfer of $1,364,000 to Science Program Direction is to consolidate all Working Capital Funds in one program.

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

	FY 2017 vs FY 2016
Mathematical, Computational, and Computer Sciences Research: With the initiation of the SC-ECP, $46,106,000 in research efforts that are on the critical path for the exascale computing initiative have been shifted to the Exascale Computing Subprogram. The computer science and computational partnerships activities will initiate efforts to investigate the impact of technologies "Beyond Moore's Law;" the SciDAC portfolio will be recompeted and expanded to initiate a partnership with BES and BER to develop new tools and technologies for the BRAIN initiative, a partnership with NNSA on seismic simulation, and partnerships supporting the Administration's Clean Energy Initiatives.	**-28,322**
High Performance Computing and Network Facilities: With the initiation of the SC-ECP, $107,894,000 of Research and Evaluation Prototypes efforts that are on the critical path for the exascale computing initiative have been shifted to the Exascale Computing Subprogram. Increased facilities funding supports: operations, including increased power costs at the facilities; initial site preparation activities for NERSC-9; LCF completion of site preparations for planned upgrades in FY 2018–19—including support for OLCF to exercise the option to take delivery of a 200pf system, which is 50pf larger than planned for in FY 2016; ALCF deployment of an interim system to help ALCF users transition to the new many-core architecture in their planned upgrade; Research and Evaluation Prototypes, which will support the Computational Sciences Graduate Fellowship at $10,000,000 and initiate testbeds for exploring computer technologies "Beyond Moore's Law" and a modest new effort in cybersecurity for HPC systems; and ESnet efforts to provide increases in bandwidth for the growing data requirements of SC facilities, such as upgrading high-traffic links to 400gbps.	**-83,498**
Exascale Computing: With the initiation of the SC-ECP, research efforts that are on the critical path for the exascale computing initiative have been shifted to the Exascale Computing subprogram.	**+154,000**
Total, Advanced Scientific Computing Research	**+42,180**

Basic and Applied R&D Coordination

Coordination across disciplines and programs is a cornerstone of the ASCR program. Partnerships within the Office of Science are mature and continue to advance the use of high performance computing and scientific networks for science. Growing areas of collaboration will be in the area of data-intensive science and readying applications for exascale. ASCR continues to have a strong partnership with NNSA for 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 (OE) and the Office of Nuclear Energy (NE), 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 (NSTC) Committee on Technology, the interagency networking and information technology R&D coordination effort, ASCR also coordinates with programs across the Federal Government. In FY 2016, cross-agency interactions and collaborations will continue, fostered by the NSCI, in coordination with OSTP.

Program Accomplishments

Advances in Materials Science to Reduce Friction in Energy Systems. Advances in materials science and nanotechnology often underpin technological improvements in energy delivery. For example, it is estimated that 1/3 of the energy consumed in passenger cars is to overcome mechanical friction.[a] In a collaboration at Argonne, involving materials scientists at the Argonne Center for Nanoscale Materials, computational scientists, and an applied-energy program laboratory fellow, the Mira supercomputer was used to identify and improve a new mechanism for reducing friction, which is feeding into the development of a hybrid material that exhibits superlubricity at the macroscale for the first time. Computer simulations revealed that when the lubricant materials—graphene and diamond-like carbon (DLC)—slid against each other, the graphene began rolling up to form hollow cylindrical "scrolls" that helped to practically eliminate friction. These so-called nanoscrolls represent a completely new mechanism for superlubricity, a state in which friction essentially disappears. Superlubricity is a highly desirable property in a host of mechanical systems. Considering that nearly one-third of every tank of fuel is spent overcoming friction in automobiles, a material that can achieve superlubricity would provide a significant economic advantage and would benefit industry and consumers alike. The research team is in the process of seeking a patent for the hybrid material, which could potentially be used for applications in dry environments, such as computer hard drives, wind turbine gears, and mechanical rotating seals for microelectromechanical and nanoelectromechanical systems. The team's groundbreaking nanoscroll discovery would not have been possible without a supercomputer like Mira. Replicating the experimental setup required simulating up to 1.2 million atoms for dry environments and up to 10 million atoms for humid environments. This work was published in Science Magazine in June 2015.

Harnessing Large Scale Turbulence to Improve Clean Coal Combustion. Fundamental understanding in chemistry and turbulence often holds the keys for improving the efficiency of large-scale combustion devices, such as those used in coal-fired power plants. Researchers at the University of Utah and its NNSA–funded Carbon Capture Multi-Disciplinary Simulation Center (CCMSC) are using the Oak Ridge LCF (OLCF) to improve modeling capabilities to enable petascale simulations to guide the design of next-generation oxy-coal boilers for clean electric energy. The scale of these boilers (up to 300 feet tall), their cost (hundreds of millions of dollars), and their ability to provide electricity for up to a million people underscore the importance of optimizing the design. The computational models developed by CCMSC use large-eddy simulation, derived from long-term investments by ASCR and Basic Energy Sciences (BES) to model turbulence in the boilers and to include ray-tracing approaches to model thermal radiation, which is the dominant mode of heat transfer. In order to show that these models are capable of being used to model a complete boiler, the team used the OLCF to develop, implement, and test Reverse Monte Carlo Ray Tracing (RMCRT), a unique technique for modeling radiative heat transfer in clean coal boilers, at large scale on both CPUs and GPUs. The OLCF GPUs reduced the team's time to solution by a factor of two and enabled simulations previously too computationally expensive to perform. These achievements lay the groundwork for full-architectural utilization—CPU and GPU—of OLCF to perform multi-physics simulation for oxy-coal boiler modeling, including combustion, fluid flow and radiative heating, necessary to achieve CCMSC's goal of simulating during the next five years clean coal boiler of between a 350MWe and 1000MWe. This technology was recently acquired by GE Power Systems.

---

[a] http://www.stle.org/resources/lubelearn/friction/

Basic Science of Converting CO_2 Into Fuel and Useful Chemicals. BES-funded basic research in the functional role of pyridinium during aqueous electrochemical reduction of CO_2 on platinum, coupled with simulations run on NERSC's Hopper supercomputer explain why submerging a platinum semiconductor into an acidic solution of pyridine and CO_2 and charging it with just 600 millivolts of electricity, the CO_2 can be transformed into formic acid, formaldehyde and methanol. Simulations reveal that, unexpectedly, the CO_2 conversion process is initiated by a reaction with hydrogen atoms bound to a platinum surface. Other mechanisms had been proposed before these simulations settled the issue. Although platinum catalysts would be too expensive to use at scale, this research opens the door to systematic exploration of catalysts that are more affordable. The findings are expected to be useful in the design and development of new technologies that can generate fuels that are consumed without producing CO_2, a long held goal in reducing the carbon footprint of chemical manufacturing and energy production.

Advancing Next Generation Nuclear Energy. Although nuclear fission of uranium has been used in commercial power plants for decades across the U.S. and generates about 20 percent of the nation's electricity in a carbon-free manner, significant fundamental intricacies of the process remain unknown. Basic research in nuclear physics, coupled with ASCR simulation capabilities are allowing scientists to delve into these questions in the study of the model system of fission of fermium-264 as it splits into two symmetric nuclei of tin-132. They combined sophisticated calculations and techniques that considered the variations in the shape of a single fermium-264 atom as it breaks apart and found that the optimal fission path is strongly impacted by nucleonic pairing, also known as superfluidity. The coupling between shape and pairing can lead to a dramatic departure from the standard picture of fission. For the first time, researchers are able to study spontaneous fission microscopically with a computational model that considers factors which were previously estimated. When extended to the more complex case of uranium fission, this research could open doors for improving the performance and safety of nuclear reactors, as well as for national security applications. The research team from the Michigan State University, funded by NNSA and using the OLCF supercomputer, is working to develop a predictive framework to describe nuclear fission.

ESnet Support of High-Energy Physics Experiments and Secretarial Recognition. In 2015, ESnet transatlantic circuit capacity was expanded to 340 gigabits/second, to support resumption of experiments at the Large-Hadron Collider at the CERN (*Organisation Européenne pour la Recherche Nucléaire*) in Geneva, Switzerland. The capacity expansion was designed and implemented on time to support the start of LHC Run2, which began in June 2015. ESnet staff received a Secretarial Honor Award—DOE's highest non-monetary recognition—for developing the On-demand Secure Circuits and Reservation System (OSCARS), a widely-used software application that creates dedicated bandwidth 'highways' for scientists to transfer massive, time-critical data sets over long distances.

Delivering on Exascale technologies. As part of the ASCR Fast Forward computer-science research activity, NVIDIA developed a signaling technology called GRS (ground-referenced signaling) that ultimately led to the development of the NV Link technology, which is a key technology in their Pascal and Volta generation GPUs. GRS is a high-speed (20 gigabits/second) low-energy (0.5-picojoule/bit) signaling technology. During the development of GRS, NVIDIA sketched a number of application scenarios including using the technology to provide a high-bandwidth channel between the CPU and GPU—and hence enabling full-bandwidth GPU access to CPU memory via NV Link. NVLink 1.0 will be available on the Pascal GP100 GPU and on the IBM Power 8+ CPU (and several other yet unannounced CPUs). It provides high-bandwidth, low-latency communication between the CPU and the GPU in a heterogeneous system, which will be critical in exascale systems. This link enables the GPU to use the full memory bandwidth of the CPU memory, which will greatly simplify the programmability of scientific and engineering algorithms on heterogeneous HPCs.

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 productivity, energy efficiency, and resiliency 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 HPC systems for data-intensive and computational science at the exascale and beyond. This requires a focus on increased parallelism, energy efficiency, and reliability.

Deriving scientific insights and knowledge from vast amounts of data flowing from Office of Science user facilities will require a focused research effort to develop the necessary theories, software 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 them.

With the initiation of the SC-ECP, research efforts that are on the critical path for the exascale computing initiative have been shifted to the Exascale Computing Subprogram.

Applied Mathematics

The Applied Mathematics activity supports the R&D 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 energy-related biological systems computationally. This activity's research underpins all of DOE's modeling and simulation efforts. Significant innovation in applied mathematics is needed to realize the potential of next generation HPC systems. High-fidelity modeling and simulation require a number of new algorithmic techniques and strategies supported by this activity, including advanced solvers for large linear and nonlinear systems of equations, time integration schemes, multi-physics coupling, methods that use asynchrony or randomness, adaptively evolving mesh techniques, algorithmic resilience, and uncertainty quantification.

Computer Science

The Computer Science activity supports research on extreme-scale computing and extreme-scale data. Information from computer vendors indicates that, because of power constraints, data movement, rather than arithmetic operations, will be a constraining factor within future systems. Memory per core is projected to decline sharply, due to power requirements and the cost of memory relative to CPUs 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 memory hardware are anticipated and will require research within this activity to develop new approaches for run-time data management and analysis.

Significant innovation in computer science is needed to realize the computational and data-analytic potential of next-generation HPC systems and other scientific user facilities in a timeframe consistent with their anticipated availability. 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 also will be significant efforts in software tools, user interfaces, the HPC software stack that can dynamically deal with time-varying energy efficiency and reliability requirements–including operating systems, file systems, compilers, and performance tools–and visualization and analytics tools that scale to extremely massive datasets. These efforts are essential to ensure DOE mission applications are able to use commercially available HPC 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. Current 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 progress at the frontiers of science is enhanced by advances in computational technology, most pressingly, the emergence of the hybrid, multi-core and many-core architectures.

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 to one another. ASCR-supported researchers helped establish critical protocols on which the internet is based. Next Generation Networking for Science research provides underpinning technologies used in international collaborations such as the Large Hadron Collider, including 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 Changes

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
Mathematical, Computational, and Computer Sciences Research $179,176,000	**$150,854,000**	**-$28,322,000**
Applied Mathematics $49,229,000	**$39,229,000**	**-$10,000,000**
Applied Mathematics continues efforts to develop new algorithmic techniques and strategies that extract scientific advances and engineering insights from massive data for DOE missions. Applied Mathematics addresses many of the challenges of exascale including: advanced solvers, uncertainty quantification, algorithmic resilience, and strategies for reducing global communications.	Applied Mathematics will continue efforts to develop new algorithmic techniques and strategies that extract scientific advances and engineering insights from massive data for DOE missions including: uncertainty quantification, algorithmic resilience, and strategies for reducing global communications.	$10,000,000 of Applied Mathematics efforts have been moved to the Exascale Computing subprogram.
Computer Science $56,848,000	**$39,296,000**	**-$17,552,000**
Computer Science continues efforts to develop software, new programming models and metrics for evaluating system status. This activity is primarily focused on addressing the challenges of exascale and data-intensive science. Emphasis remains on efforts to promote ease of use; increase parallelism, energy efficiency, and reliability, and ensures that research efforts are tightly coupled to application requirements and developments in industry, particularly those identified by the co-design centers and developed in partnerships supported by the Research and Evaluation Prototypes activity.	Computer Science will continue efforts to develop software, new programming models, new operating systems and efforts to promote ease of use. As Moore's law reaches its final phase, impacts will be seen in several technology areas, including HPC and knowledge extraction from large datasets resulting from next-generation scientific experiments. Given the Department's significant investment in these areas, it is essential that our research community understands and is ready to mitigate these impacts to our mission applications.	$20,106,000 of Computer Science efforts have been moved to the Exascale Computing subprogram. The computer science activity will initiate new efforts to begin to investigate the impact of technologies "Beyond Moore's Law."

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
Computational Partnerships $47,918,000	**$45,596,000**	**-$2,322,000**
The SciDAC Institutes are recompeted at the end of FY 2016. These Institutes continue to provide the bridge between the core research program and the DOE science applications. The development of SciDAC tools and resources by the Institutes is primarily for use on computational systems, such as those existing and planned for at the Oak Ridge and Argonne Leadership Computing Facilities, the National Energy Research Scientific Computing Center, and similar world-class computing facilities over the following five years. In addition, the exascale Co-design centers will undergo a comprehensive external peer review in FY 2015 to document progress and impact, and to inform the recompetition of these efforts in FY 2016.	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 science application partnerships, with the other Office of Science programs, will be recompeted and expanded in FY 2017 to address the challenges in moving toward larger data sets and more complex computing systems on the path to exascale. FY 2017 marks the beginning of the fourth iteration of the SciDAC portfolio. This highly successful program has never been more important to bridge the gap between core research efforts in computer science and applied math and the applications supported by the other Office of Science programs. The focus on SciDAC over the next four years will be on readying Office of Science applications to harness the potential of the upgraded ASCR leadership and production computing facilities with priority emphasis on efforts that are needed to advance the science goals in partnerships with the Offices of Basic Energy Sciences (BES), Biological and Environmental Research (BER), High Energy Physics, Nuclear Physics, and Fusion Energy Sciences (FES). ASCR will continue to work with the DOE Applied Energy programs and other Federal agencies, in support of the whole-of-government objective of the NSCI. As part of ASCR efforts to understand the impacts of technologies “Beyond Moore’s Law,” this activity will initiate partnerships within the Office of Science for application specific efforts to begin to explore these technologies.	$16,000,000 of co-design efforts have been moved to the Exascale Computing subprogram. Funding supports recompetition and expansion of the SciDAC partnerships including new efforts with BES, BER and FES supporting the Administration’s Clean Energy Initiatives. Funding also initiates a partnership with BES and BER to develop new tools and technologies for the BRAIN initiative, a partnership with NNSA on seismic simulation, and partnerships within the Office of Science to begin to explore technologies “Beyond Moore’s Law.”

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
Next Generation Networking for Science $19,000,000	**$19,000,000**	**$0**
The Next Generation Networking for Science activity continues to work closely with SC user facilities and applications, to develop the necessary tools—networking software, middleware and hardware—to address the challenges of moving, sharing and validating massive quantities of data via next generation optical networking technologies. This focus allows DOE scientists to productively collaborate regardless of the geographical distance between scientists and user facilities or the size of the data.	The Next Generation Networking for Science activity will continue to work closely with SC user facilities and applications, to develop the necessary tools—networking software, middleware and hardware—to address the challenges of moving, sharing and validating massive quantities of data via next generation optical networking technologies.	No change
SBIR/STTR $6,181,000	**$7,733,000**	**+$1,552,000**
In FY 2016, SBIR/STTR funding is set at 3.45% of non-capital funding.	In FY 2017, SBIR/STTR funding is set at 3.65% of 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 Office of Science research facilities, generate many petabytes of data each year. Moving data to where it is needed requires advanced scientific networks and related technologies provided through High Performance Network Facilities and Testbeds, which includes the Energy Science Network (ESnet). Finally, the operation of the facilities also includes investments to ensure the facilities remain state-of-the-art and can accept future systems such as electrical and mechanical system enhancements.

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. The Research and Evaluation Prototypes activity also prepares researchers to effectively use the next generation of scientific computers and seeks to reduce risk for future major procurements.

Allocation of computer time at ASCR facilities follows the peer-reviewed, public-access model used by other SC scientific user facilities. To help address the workforce issues at the ASCR facilities, each facility established a postdoctoral training program in FY 2015 for high-end computational science and engineering. These programs teach PhD scientists with limited experience in HPC the skills to be computational scientists adept at using high performance production and leadership systems.

High Performance Production Computing

This activity supports NERSC, which delivers high-end production computing services for the SC research community. Approximately 6,000 computational scientists in about 800 projects use NERSC annually to perform scientific research across a wide range of disciplines including astrophysics, chemistry, climate modeling, materials, high energy and nuclear physics, fusion, 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 Office of Science computing facilities. In FY 2015, NERSC moved into the new Computational Research and Theory building located on the Lawrence Berkeley National Laboratory campus.

NERSC 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. This gap between demand and capacity exists despite upgrades to the primary computing systems approximately every three years. NERSC regularly gathers requirements from Office of Science domain programs through a long-established, robust process and uses these requirements to inform 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 exascale computing.

Leadership Computing Facilities

The LCFs enable open scientific applications, including industry applications, to harness the potential of leadership computing to advance science and engineering. The success of this effort is built on the gains made in Research and Evaluation Prototypes and ASCR research efforts. Another LCF strength is the staff, which operate and maintain the forefront computing resources and provide support to Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects, ASCR Leadership Computing Challenge (ALCC) 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's (OLCF) 27 petaflop (pf) system is one of the most powerful computers in the world for scientific research, and is ranked number two on the June 2015 Top 500 list.[a] Through allocations at the OLCF, many applications, including the water cycle and cryosphere systems in advanced climate simulations, developing models of astrophysical explosions, direct-numerical simulation of turbulent combustion flexible-fuel gas turbines, dynamical simulations of magnetic fields in high-energy-density plasmas, first-principles based statistical physics of alloys and functional materials, simulations of high-temperature superconductors, molecular design of next-generation nanostructured polymer electrolytes, simulation of fundamental energy conversion processes in cells, simulations of neutron transport in fast-fission reactor cores, and earthquake simulations, are running at the multi-petaflop scale. OLCF staff shares its expertise with industry to broaden the benefits of petascale computing for the Nation. For example, OLCF works with industry to reduce the need for costly physical prototypes and physical tests in the development of high-technology products.

The Argonne Leadership Computing Facility (ALCF) operates a 10-pf IBM Blue Gene Q (Mira), developed through a joint research project with support from the NNSA, industry, and ASCR's Research and Evaluation Prototypes activity. This HPC system achieves high performance with relatively lower electrical power consumption than other current petascale computers. The ALCF will also begin operations of an 8.5 pf Intel-based machine to prepare their users for the ALCF-3 upgrade in 2019.

The ALCF and OLCF systems are architecturally distinct, consistent with DOE's strategy to foster a diversity of capabilities that provides the Nation's HPC user community the most effective resources. 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 scientists and engineers to understand the fundamental physics of turbulent mixing to transform product design and to achieve improved performance, lifespan and efficiency.

The demand for 2015 INCITE allocations at the LCFs outpaced the available resources by a factor of three.

Research and Evaluation Prototypes

Research and Evaluation Prototypes (REP) has a long history of partnering with U.S. vendors to develop future computing technologies and testbeds that push the state-of-the-art and allowed DOE researchers to better understand the challenges and capabilities of emerging technologies. This activity supports testbeds for next-generation systems and for technologies "Beyond Moore's Law." This activity also supports some near-term research efforts needed by the ASCR facilities and focused on their unique needs—such as cybersecurity efforts focused on the unique challenges of open HPC systems.

In addition, this activity partners with the NNSA on the Computational Sciences Graduate Fellowship (CSGF).

With the initiation of the SC-ECP, research and evaluation prototype efforts that are on the critical path for the exascale computing initiative and activities, such as the vendor partnerships on critical technologies, nodes and system integration, have been shifted to the Exascale Computing Subprogram.

High Performance Network Facilities and Testbeds

The Energy Sciences Network (ESnet) provides the national and international network and networking infrastructure connecting DOE science facilities, experiments, 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 400-gbps technologies and software-defined networking services that will be required to keep pace with the expected data volume.

---

[a] http://www.top500.org/lists/2014/11/

Advanced Scientific Computing Research
High Performance Computing and Network Facilities

Activities and Explanation of Changes

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
High Performance Computing and Network Facilities $441,824,000	**$358,326,000**	**-$83,498,000**
High Performance Production Computing $86,000,000	**$92,145,000**	**+$6,145,000**
Supports installation, acceptance and operation of the NERSC high-end capability systems (NERSC-7 and NERSC-8) including increased power costs, lease payments, and user support and continuation of the post-doctoral training program for high-end computational science and engineering.	In January 2017, NERSC will begin production operations of the NERSC-8 system. Named “Cori” after Nobel Laureate Gerty Cori, this system is an approximately 30-pf Cray XC supercomputer with Intel Xeon Phi processors. The full production system will provide roughly four times the CY 2016 NERSC capacity. Demand for production computing for the Office of Science continues to grow along with system capability and the rapid increase in data from experiments. To help NERSC keep pace, the Department approved the NERSC-9 Mission Need Statement in August 2015. An RFP was issued in FY 2016 and the vendor will be selected in FY 2017 to deliver a pre-exascale system, with three to five times the capacity of NERSC-8, in 2020.	Increase supports production operations of NERSC-8, including increased power costs, lease payments, and user support, and the start of site preparations for NERSC-9 .
Leadership Computing Facilities $181,317,000	**$187,000,000**	**+$5,683,000**
Support operation and allocation of the 27-pf Titan system at the OLCF and the 10-pf Mira system at the ALCF through INCITE and ALCC. This includes lease payments, power, and user support. Also supports preparations—such as power, cooling and cabling at the LCFs to support 75-200 pf upgrades at each facility and continuation of the post-doctoral training program for high-end computational science and engineering.	Site preparations for planned upgrades at both LCF facilities will be completed in 2017. The OLCF will begin to install cabinets for the new 200-pf IBM Power9 heterogeneous supercomputer with NVIDIA Volta GPUs at the end of FY 2017. Installation, testing, early science access and transition to operations for the system, named Summit, are all planned for FY 2018. This upgrade will provide approximately five times the capability of Titan.	Increase supports completion of site preparations for planned upgrades at both LCF facilities. These include finalizing power and cooling capacity, as well as space and weight requirements because the final specification of the machines is now known for both facilities. Increase also supports installation and operations of an interim Intel-Cray computing system at the ALCF to transition users to the new many-core architecture

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
	The ALCF will complete site preparations in early FY 2018 for its planned upgrade in FY 2019 to an Intel-Cray supercomputer, called Aurora, to be built with 3rd Generation Intel Xeon Phi many-core processors. This system is a dramatic change from the Current Mira system, which is an IBM BlueGene Q supercomputer. Therefore, the ALCF will install and operate an 8.5-pf interim system, called Theta, in early FY 2017 based on Intel's second-generation Xeon Phi processor. This system will serve as an early production system to transition ALCF users to the new architecture. The LCF upgrades will advance energy and manufacturing technologies, as well as our fundamental understanding of the universe, while maintaining the United States' global leadership in HPC on the path to exascale computing.	and exercise the option to acquire a 200-pf machine at the OLCF expanding the current plan of record by 50 pf.
Leadership Computing Facility at ANL: $77,000,000	$80,000,000	+$3,000,000
Leadership Computing Facility at ORNL: $104,317,000	$107,000,000	+$2,683,000
Research and Evaluation Prototypes $121,471,000	**$17,890,000**	**-$103,581,000**
REP supports efforts to improve the energy efficiency and reliability of critical technologies such as memory, processors, network interfaces and interconnects. The compute node is the basic building block of a high performance computer and all of these technologies come together in the node. Therefore, REP will competitively select R&D partnerships with U.S. vendors to initiate the design and development of node and system designs suitable for exascale systems. These efforts will influence the development of prototypes that advance DOE goals and are based on the results of the *Fast Forward and Design Forward* investments. This is an essential component of the Department's exascale computing plan and a key step in the vendor's productization efforts.	Availability of experienced and knowledgeable workforce issues continues to be of vital importance to the ASCR facilities and to the NSCI exascale goals. The CSGF program plays an important role in providing future DOE leaders in HPC and computational science. Therefore, Research and Evaluation Prototypes will continue support for the program at $10,000,000 in FY 2017. This activity will also support a modest effort in cybersecurity research that is focused on the unique challenges of open-science HPC systems and a small scale testbed for researchers explore technologies "Beyond Moore's Law".	With the initiation of the SC-ECP, REP funding of $107,894,000 (a reduction of $3,577,000 from the FY 2016 level for these efforts) has been shifted to the newly established Exascale Computing subprogram. Increase is for a new effort in cybersecurity and a testbed for "Beyond Moore's Law" research efforts.

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
Support is provided for non-recurring engineering efforts in support of ASCR facilities. To emphasize the vital importance of the CSGF program to the ASCR facilities and to our exascale goals, Research and Evaluation Prototypes supports the program at $10,000,000.		
High Performance Network Facilities and Testbeds $38,000,000	**$45,000,000**	**+$7,000,000**
ESnet operates the national and international network infrastructure to support critical DOE science applications, SC facilities and scientific collaborations around the world through 100 gbps production network and begin upgrade to 400-gbps testbed for networking testing and research.	ESnet has become a critical enabler of large-scale, data-intensive science in the U.S. The volume of data transferred by ESnet is growing roughly 66% per year, twice the rate of the commercial Internet. This request provides bandwidth increases to support data requirements of all Office of Science facilities; upgrades selected high-traffic links to 400-gbps technology; supports transatlantic access to LHC data; and extends direct science engagement efforts to improve end-to-end network performance between DOE facilities and U.S. universities. The request also continues ongoing enhancement of network architectures and tools now widely deployed through the DOE and university systems in the US: Science DMZ, perfSONAR, Data Transfer Nodes, and OSCARS. Additionally, the request supports applied R&D necessary to maintain ESnet's status as a world-leading scientific research network, and to support a network testbed focused on prototyping and operationalizing future network architectures such as Software-Defined Networking and Named-Data Networking.	Increase supports operations and staff for the ESnet, including upgrading high-traffic links to 400 gbps; transatlantic access to LHC data; expansion of science engagement efforts; and research and tool development.
SBIR/STTR $15,036,000	**$16,291,000**	**+$1,255,000**
In FY 2016, SBIR/STTR funding is set at 3.45% of non-capital funding.	In FY 2017, SBIR/STTR funding is set at 3.65% of non-capital funding.	

Advanced Scientific Computing Research
Exascale Computing

Description

The Office of Science Exascale Computing Project (SC-ECP) in the Exascale Computing subprogram captures ASCR's participation in the U. S. Department of Energy's ECI, to ensure the hardware and software R&D, including applications software, for a capable exascale system is completed in time to meet the scientific and national security mission needs of the mid-2020s.

On July 29, 2015 President Obama established, by Executive Order, the National Strategic Computing Initiative (NSCI) to maximize the benefits of High Performance Computing (HPC) for U.S. economic competitiveness and scientific discovery. DOE is one of the three lead agencies for the NSCI and is specifically assigned the responsibility for executing a program, joint between Office of Science and NNSA, to develop a capable exascale computing program with an emphasis on sustained performance on relevant applications and analytic computing to support DOE missions.

The SC-ECP comprises R&D and delivery of exascale computers and is not a traditional construction project. The SC-ECP will be managed following the principles of DOE Order 413.3B, *Program and Project Management for the Acquisition of Capital Assets*, which has been used by the Office of Science for the planning, design, and construction of all of its major projects, including the LCFs at Argonne and Oak Ridge National Laboratories and NERSC at Lawrence Berkeley National Laboratory. Computer acquisitions will use a tailored version of Order 413.3B. The first four years of SC-ECP will focus on research in software (new algorithms and methods to support application and system software development) and hardware (node and system design). During the last six years of the SC-ECP, activities will focus on delivering application software, the system software stack, and hardware technologies that will be deployed in the exascale systems.

FY 2017 funding will support codesign activities with a representative subset of mission applications; efforts to develop an exascale software stack, new programming models and efforts to promote ease of use; increase parallelism, energy efficiency, and reliability; and the initiation of vendor partnerships to design and develop scalable prototypes of exascale systems.

Activities and Explanation of Changes

FY 2016 Enacted	FY 2017 Request	Explanation of Changes FY 2017 vs FY 2016
17-SC-20 Office of Science Exascale Computing Project (SC-ECP) $0	**$154,000,000**	**+$154,000,000**
Research efforts that are on the critical path for the ECI have previously been funded within Applied Mathematics, Computer Sciences, Computational Partnerships, and Research and Evaluation Prototypes activities.	FY 2017 funding will support codesign activities with a representative subset of mission applications; efforts to develop an exascale software stack, new programming models and efforts to promote ease of use; increase parallelism, energy efficiency, and reliability and the initiation of vendor partnerships to design and develop scalable prototypes of exascale systems.	Increase reflects the transfer of exascale research efforts previously funded within Applied Mathematics, Computer Sciences, Computational Partnerships, and Research and Evaluation Prototypes activities.

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. The following table shows the targets for FY 2015 through 2017. Details on the Annual Performance Report can be found at http://energy.gov/cfo/reports/annual-performance-reports.

	FY 2015	FY 2016	FY 2017
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	Conduct an external peer review of the three original co-design centers to document progress, impact, and lessons learned.	Fund two teams to develop exascale node designs.	Fund two teams to develop programming environments for exascale computing systems.
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 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Request	FY 2017 vs FY 2016
Capital operating expenses							
Capital equipment	n/a	n/a	8,000	13,100	6,000	5,000	-1,000

Funding Summary ($K)

	FY 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Request	FY 2017 vs FY 2016
Research	228,169	221,018	294,466	315,011	+20,545
Scientific user facility operations	295,242	302,393	305,317	324,145	+18,828
Other	17,589	0	21,217	24,024	+2,807
Total, Advanced Scientific Computing Research	**541,000**	**523,411**	**621,000**	**663,180**	**+42,180**

Advanced Scientific Computing Research
Scientific User Facility Operations ($K)

The treatment of user facilities is distinguished between two types: TYPE A facilities that offer users resources dependent on a single, large-scale machine; TYPE B facilities that offer users a suite of resources that is not dependent on a single, large-scale machine.

Definitions:
Achieved Operating Hours – The amount of time (in hours) the facility was available for users.

Planned Operating Hours –

- For Past Fiscal Year (PY), the amount of time (in hours) the facility was planned to be available for users.
- For Current Fiscal Year (CY), the amount of time (in hours) the facility is planned to be available for users.
- For the Budget Fiscal Year (BY), based on the proposed budget request the amount of time (in hours) the facility is anticipated to be available for users.

Optimal Hours – The amount of time (in hours) a facility would be available to satisfy the needs of the user community if unconstrained by funding levels.

Percent of Optimal Hours – An indication of utilization effectiveness in the context of available funding; it is not a direct indication of scientific or facility productivity.

- For BY and CY, Planned Operating Hours divided by Optimal Hours expressed as a percentage.
- For PY, Achieved Operating Hours divided by Optimal Hours.

Unscheduled Downtime Hours - The amount of time (in hours) the facility was unavailable to users due to unscheduled events. NOTE: For type "A" facilities, zero Unscheduled Downtime Hours indicates Achieved Operating Hours equals Planned Operating Hours.

	FY 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Request	FY 2017 vs FY 2016
TYPE A FACILITIES					
NERSC	**$75,605**	**$75,905**	**$86,000**	**$92,145**	**+6,145**
Number of Users	5,608	5,608	5,608	6,000	+392
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	8,585	8,585	8,585	8,322[a]	-263
Optimal hours	8,585	8,585	8,585	8,322	-263
Percent optimal hours	N/A	N/A	N/A	N/A	N/A
Unscheduled downtime hours	N/A	N/A	N/A	N/A	N/A

---

[a] Due to the planned upgrade NERSC will schedule less hours for FY 2017. However, the significant increase in the capacity of this upgrade system will allow NERSC to deliver an increase in computing time for users despite the reduced schedule.

	FY 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Request	FY 2017 vs FY 2016
OLCF	**$104,317**	**$108,902**	**$104,317**	**$107,000**	**+$2,683**
Number of Users	1,064	1,064	1,064	1,064	-
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	7,008	7,008	7,008	7,008	-
Optimal hours	7,008	7,008	7,008	7,008	-
Percent optimal hours	N/A	N/A	N/A	N/A	N/A
Unscheduled downtime hours	N/A	N/A	N/A	N/A	N/A
ALCF	**$80,320**	**$81,796**	**$77,000**	**$80,000**	**+$3,000**
Number of Users	1,434	1,434	1,434	1,434	-
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	7,008	7,008	7,008	7,008	-
Optimal hours	7,008	7,008	7,008	7,008	-
Percent optimal hours	N/A	N/A	N/A	N/A	N/A
Unscheduled downtime hours	N/A	N/A	N/A	N/A	N/A
ESnet	**$35,000**	**$35,790**	**$38,000**	**$45,000**	**+$7,000**
Number of users[a]	N/A	N/A	N/A	N/A	N/A
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	8,760	8,760	8,760	8,760	-
Optimal hours	8,760	8,760	8,760	8,760	-
Percent optimal hours	N/A	N/A	N/A	N/A	N/A
Unscheduled downtime hours	N/A	N/A	N/A	N/A	N/A

---

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

	FY 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Request	FY 2017 vs FY 2016
Total Facilities	**$295,242**	**$302,393**	**$305,317**	**$324,145**	**+$18,828**
Number of Users[a]	8,106	8,106	8,106	8,498	+392
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	31,361	31,361	31,361	31,098	-263
Optimal hours	31,361	31,361	31,361	31,098	-263
Percent of optimal hours[b]	N/A	N/A	N/A	N/A	N/A
Unscheduled downtime hours	N/A	N/A	N/A	N/A	N/A

Scientific Employment

	FY 2015 Enacted	FY 2015 Current	FY 2016 Enacted	FY 2017 Estimate	FY 2017 vs FY 2016
Number of permanent Ph.D.'s (FTEs)	548	548	584	636	+52
Number of postdoctoral associates (FTEs)	137	137	146	165	+19
Number of graduate students (FTEs)	428	428	460	486	+26
Other scientific employment (FTEs)[c]	234	234	247	257	+10

---

[a] Total users only for NERSC, OLCF, and ALCF.

[b] For total facilities only, this is a "funding weighted" calculation FOR ONLY TYPE A facilities: $\frac{\sum_1^n[(\%OH\ for\ facility\ n)\times(funding\ for\ facility\ n\ operations)]}{Total\ funding\ for\ all\ facility\ operations}$

[c] Includes technicians, engineers, computer professionals and other support staff.

17-SC-20 Office of Science Exascale Computing Project (SC-ECP)

1. Significant Changes and Summary

Significant Changes

In FY 2016, the President's Budget Request included funding to initiate research, development, and computer-system procurements to deliver an exascale (10^{18} operations per second) computing capability by the mid-2020s. This activity, referred to as the Exascale Computing Initiative (ECI), is in partnership with the National Nuclear Security Administration (NNSA) and addresses DOE's science and national security mission requirements. In FY 2016, an Exascale Crosscut aggregated Office of Science programs (specifically the Advanced Scientific Computing Research (ASCR), Basic Energy Sciences (BES), and Biological and Environmental Research (BER) programs) and NNSA exascale activities.

In FY 2017, the Office of Science component of the ECI is partitioned into the Office of Science Exascale Computing Project (SC-ECP) within a new Exascale Computing subprogram in ASCR, and includes only those activities required for the delivery of exascale computers. The FY 2017 Request for this project is $154,000,000, which represents only those activities, and is a decrease of $3,894,000 below the FY 2016 Enacted level of $157,894,000[a]. Other activities related to the ECI but outside of the scope of the delivery of exascale computers are not within the SC-ECP, though they do remain in the scope of the ECI and are funded through the Exascale Computing subprogram in ASCR, and through subprograms within BES and BER. This Project Data Sheet is for the SC-ECP only; prior-year activities related to the SC-ECP are also included.

Summary

On July 29, 2015, President Obama established by Executive Order the National Strategic Computing Initiative (NSCI) to maximize the benefits of High Performance Computing (HPC) for U.S. economic competitiveness, scientific discovery, and national security. DOE is one of three lead federal agencies for the NSCI, and it is specifically assigned the responsibility for executing the ECI. As a lead agency, DOE will work with other agencies identified in the NSCI to implement the objectives of the NSCI and to address the wide variety of needs across the Federal Government.

In FY 2017, SC-ECP funding will support project management; development of project documentation; conduct of co-design activities with a representative subset of mission applications; research and development of exascale systems software and tools needed for exascale programming; and vendor partnerships. The estimated Total Project cost range of the SC-ECP is $1.7 billion to $2.7 billion.

2. Critical Milestone History

(fiscal quarter or date)

	CD-0	Conceptual Design Complete	CD-1/3A	CD-2	Final Design Complete	CD-3B	D&D Complete	CD-4
FY 2017	3Q FY 2016	TBD	TBD	TBD	TBD	TBD	N/A	TBD

CD-0 – Approve Mission Need
CD-1 – Approve Alternative Selection and Cost Range
CD-2 – Approve Performance Baseline
CD-3A – Approve phase one funding of hardware and software research projects and application development.
CD-3B – Approve phase two funding of hardware and software development, and exascale system contract options.
CD-4 – Approve Project Completion

3. Project Cost History

The preliminary cost range for the SC-ECP is estimated to be between $1.7 billion and $2.7 billion. The cost range will be updated at CD-0 and CD-1, and a project baseline (scope, schedule, and cost) will be established at CD-2.

---

[a] In FY 2016, ASCR's ECI activities were not partitioned into a separate program/project.

4. Project Scope and Justification

Scope

Four well-known challenges[a] determine the requirements of the SC-ECP. These challenges are:

- *Parallelism:* Systems must exploit the extreme levels of parallelism that will be incorporated in an exascale computer;
- *Resilience:* Systems must be resilient to permanent and transient faults;
- *Energy Consumption:* System power requirements must be no greater than 20-30 MW; and
- *Memory and Storage Challenge*: Memory and storage architectures must be able to access and store information at anticipated computational rates.

The realization of an exascale system that addresses parallelism, resilience, energy consumption, and memory/storage will involve tradeoffs among hardware (processors, memory, energy efficiency, reliability, interconnectivity); software (programming models, scalability, data management, productivity); and algorithms. To address this, the scope of the SC-ECP has four focus areas:

1. *Hardware Technology:* The Hardware Technology focus area supports vendor-based research and development activities required to deploy at least two exascale systems with diverse architectural features. Within this focus area, a node design effort targets component technologies needed to build exascale nodes, including the required software, while a system design effort performs the engineering and R&D activities required to build a full exascale computer and the required systems software.
2. *System Software Technology:* The System Software Technology focus area spans low-level operational software to programming environments for high-level applications software development, including the software infrastructure to support large data management and data science for the DOE at exascale.
3. *Application Development:* The Application Development focus area includes: extreme parallelism, reliability and resiliency, deep hierarchies of hardware processors and memory, scaling to larger systems, and data-intensive science.
4. *Exascale Systems:* The Exascale Systems focus area supports advanced system engineering development by the vendors needed to produce capable exascale systems. System procurement activities will be executed in coordination with each DOE HPC facility's existing system acquisition timelines. It also includes acquisition and support of prototypes and testbeds for the application, software, and hardware testing activities. No civil construction is within the scope of the SC-ECP.

The SC-ECP will be managed following the principles of DOE Order 413.3B, *Program and Project Management for the Acquisition of Capital Assets*, which has been used by the Office of Science for the planning, design, and construction of all of its major projects, including the Leadership Computing Facilities at Argonne and Oak Ridge National Laboratories and NERSC at Lawrence Berkeley National Laboratory. Computer acquisitions use a tailored version of Order 413.3B. The first four years of SC-ECP will be focused on research in software (new algorithms and methods to support application and system software development) and hardware (node and system design) and these costs will be reported as Other Project Costs. During the last five years of the project, project activities will focus on delivering application software, the system software stack, and hardware technologies that will be deployed in the exascale systems and these costs will be included in the Total Estimated Costs for the project.

---

[a] http://www.isgtw.org/feature/opinion-challenges-exascale-computing

5. Preliminary Financial Schedule

	(dollars in thousands)		
	Appropriations	Obligations	Costs
Total Estimated Cost (TEC) (Delivery of Applications Development, System Software Technology Hardware Technology and Exascale Systems)			
FY 2016[a]	0	0	0
FY 2017	0	0	0
FY 2018[b] – FY 2025	TBD	TBD	TBD
Subtotal	TBD	TBD	TBD
Total, TEC	TBD	TBD	TBD
Other project costs (OPC) (Research for Application Development , System Software Technology and Hardware Technology)			
FY 2016[a]	157,894	157,894	87,370
FY 2017	154,000	154,000	156,000
FY 2018	TBD	TBD	68,524
FY 2019 – FY 2026	TBD	TBD	TBD
Subtotal	TBD	TBD	TBD
Total, OPC	TBD	TBD	TBD
Total Project Costs (TPC)			
FY 2016[a]	157,894	157,894	87,370
FY 2017	154,000	154,000	156,000
FY 2018[b]	TBD	TBD	68,524
FY 2019 – FY 2026	TBD	TBD	TBD
Subtotal	TBD	TBD	TBD
Total, TPC	TBD	TBD	TBD

---

[a] Funding was provided to ASCR in FY 2016 to support the Department's ECI efforts. For completeness, that information is shown here.

[b] The project is currently pre-CD-0 so yearly planning numbers are not available.

6. Details of the 2017 Project Cost Estimate

The SC-ECP will be baselined at CD-2. The current cost estimate is based on pre-CD-0 information. The estimated Total Project Cost for the SC-ECP is represented in the table below.

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Production Ready Software	TBD	N/A	N/A
Hardware Build	TBD	N/A	N/A
Total, TEC	TBD	N/A	N/A
Other Project Costs (OPC) (Research)			
Planning/Project Mgmt	8,000	N/A	N/A
Application Development	85,000	N/A	N/A
Software Research	87,000	N/A	N/A
Hardware Research	131,894	N/A	N/A
Total OPC	TBD	N/A	N/A
Total, TPC	TBD	N/A	N/A

7. Schedule of Appropriation Requests

($K)

Request Year		FY 2016[a]	FY 2017	FY 2018	FY 2019	FY 2020	FY 2021	Outyears	Total
FY 2017	TEC	0	0	TBD	TBD	TBD	TBD	TBD	TBD
	OPC	157,894	154,000	TBD	TBD	TBD	TBD	TBD	TBD
	TPC	157,894	154,000	TBD	TBD	TBD	TBD	TBD	TBD

8. Related Operations and Maintenance Funding Requirements

System procurement activities for the capable exascale computers are not part of the SC-ECP. The exascale computers will become part of existing facilities and operations and maintenance funds and will be included in the ASCR facilities' operations budget.

Start of Operation	2024
Expected Useful Life (number of years)	5
Expected Future start of D&D for new construction (fiscal quarter)	4Q 2030

9. D&D Funding Requirements

N/A, no construction.

---

[a] Funding was provided to ASCR in FY 2016 to support the Department's ECI efforts. For completeness, that information is shown here.

10. Acquisition Approach

The early years of the SC-ECP, approximately four years in duration, will support R&D directed at achieving system performance targets for parallelism, resilience, energy consumption, and memory and storage. The second phase of approximately five years duration will support finalizing applications and system software, the procurement of an exascale computer system, and start of operations.