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 and engineering, in partnership with the research community, including U.S. industry. The ASCR program gives the science and technology community, including U.S. industry, access to world-class supercomputers and the tools to use them for science and engineering. ASCR accomplishes this by developing and maintaining world-class computing and network facilities for science; and advancing research in applied mathematics, computer science, and advanced networking.

For over half a century, the U.S. maintained world-leading computing capabilities through sustained investments in research and the development and deployment of new computing systems. The benefits of U.S. computational leadership have been enormous – huge gains in workforce productivity, an acceleration of progress in both science and engineering, advanced manufacturing techniques and rapid prototyping, stockpile stewardship without testing, and the ability to explore, understand and harness natural and engineered systems that are too large, too complex, too dangerous, too small or too fleeting to explore experimentally. As the Council on Competitiveness noted and documented in a series of case studies, "A country that wishes to out-compete in any market must also be able to out-compute its rivals."[a]

The U.S. is now entering an era where advances in computing capabilities are becoming increasingly costly and risky, while at the same time, U.S. dominance in computing is under threat from significant new investments in Asia and Europe. We cannot afford to fall behind in an area that impacts every sector of our economy and every field of science and engineering. Therefore, this Budget Request increases our investments in the Exascale Computing Initiative (ECI) to accelerate the project and intends to accelerate delivery of at least one exascale-capable system in 2021.

The Department of Energy (DOE) and its predecessor organizations have long played a key role in advancing U.S. computing capabilities in partnership with U.S. computing vendors and researchers. Computing is a fast-paced industry, but sustained progress depends upon significant gains in numerous areas of fundamental research including: advanced lithography, nano-scale materials science, applied mathematics and computer science – areas where DOE has provided long-term investments and world-leading capabilities. Because DOE partners with High Performance Computing (HPC) vendors to accelerate and influence 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. Public-private partnership remains vital as we push our state-of-the-art fabrication techniques to their limit to develop an exascale-capable (a billion billion operations per second) system while simultaneously preparing for what follows the end the current technology.

Maximizing the benefits of U.S. leadership computing in the coming decades will require an effective national response to increasing demands for computing capabilities and performance, emerging technological challenges and opportunities, and competition with other nations. As one of the leading Federal agencies, DOE will sustain and enhance its support for 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 partnerships, 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 broadly collaborating with industry and academia.
- Develop a comprehensive technical and scientific approach to rapidly transition research on hardware, system software, development tools, and applications into production.

Within the context of this coordinated Federal strategy, the DOE Office of Science (SC) and the DOE National Nuclear Security Administration (NNSA) are overseeing an ECI. The ECI, which began in FY 2016, is a partnership between SC and

---

[a] Final report from the High Performance Computing Users Conference: Supercharging U.S. Innovation & Competitiveness, held in July 2004.

NNSA to perform the research and development (R&D) necessary to overcome key exascale challenges in parallelism, energy efficiency, and reliability, leading to intended deployment of exascale systems in 2021. The ECI's goal for an exascale-capable system is a fifty-fold increase in sustained performance over today's computing capabilities, with applications that address next-generation science, engineering, and data problems. The ECI focuses on delivering advanced simulation through an exascale-capable 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.

The SC FY 2018 Budget Request funds two components of the ECI: planning, site preparations, and non-recurring engineering at the Leadership Computing Facilities (LCF) to prepare for intended deployment of at least one exascale system in 2021, and the ASCR-supported Office of Science Exascale Computing Project (SC-ECP), proposed in the FY 2017 Request and again in this FY 2018 Request, which includes only the R&D activities required to develop exascale-capable computers.

The scope of the SC-ECP has three 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; and
- *Application Development:* The Application Development focus area includes: working with scientific application areas to address the challenges of extreme parallelism, reliability and resiliency, deep hierarchies of hardware processors and memory, scaling to larger systems, and data-intensive science.

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 SC for the planning, design, and construction of all of its major projects, but tailored to address the challenges of this fast-paced, research focused, public-private, HPC project.

Overall project management for SC-ECP will be conducted via a Project Office that has been established at Oak Ridge National Laboratory (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 Memorandum of Agreement has been signed between the six DOE national laboratories participating in SC-ECP: Lawrence Berkeley National Laboratory (LBNL), ORNL, Argonne National Laboratory (ANL), Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL), and the Project Office is now executing the project and coordinating among partners.

Highlights of the FY 2018 Budget Request

The FY 2018 Budget Request for ASCR increases our investments in the ECI with the intention to accelerate the project and deliver at least one exascale-capable system in 2021. To ensure continued progress during and after the exascale project, this Request also maintains support for ASCR's fundamental research in Applied Mathematics, Computer Science, and Computational Partnerships. Funding for the ASCR facilities supports continued operations of the upgraded Oak Ridge Leadership Computing Facility (OLCF), the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing Center (NERSC), and the Energy Sciences Network (ESnet). Funding also supports initiation of an upgrade of the ESnet and site preparations at the LCFs in support of an intended 2021 delivery of at least one exascale computing system.

<u>Mathematical, Computational, and Computer Sciences Research</u>

Recognizing that Moore's Law (microchip feature sizes reduces by a factor of two approximately every two years) is nearing an end due to limits imposed by fundamental physics, ASCR began new activities in FY 2017 to explore future computing, such as quantum information, that are not based on silicon microelectronics. In FY 2018, ASCR will continue the research and computational partnerships with the Basic Energy Sciences (BES), Biological and Environmental Research (BER) and High Energy Physics (HEP) programs aimed at understanding the challenges that quantum information and neuromorphic 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 emerging technologies.

Activities in Applied Mathematics and Computer Science provide the foundation for increasing the capability of the national HPC ecosystem. In FY 2018, these activities will continue to develop the methods, software, and tools to ensure DOE applications can fully exploit the most advanced computing systems available today 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, which are being recompeted and expanded in FY 2017. This allows the other scientific programs in SC to more effectively use the current and immediate next generation HPC facilities. The focus of the SciDAC portfolio will continue to be on developing the mission critical applications of the other SC programs. These efforts will be informed by the research results emerging from the ECI and will, whenever possible, incorporate the software, methods, and tools developed by that initiative.

The Next Generation Networking for Science activity will be eliminated. Collaboratory efforts, currently supported by this activity, will be supported by computational partnerships to strengthen the interconnectivity of these efforts. Networking R&D will be supported by the High Performance Network Facilities and Testbeds activity and will focus on the unique needs of the ESnet in support of a significant upgrade of the ESnet to address the increased data flows from the exascale systems and other SC user facilities.

<u>High Performance Computing and Network Facilities</u>
In FY 2018, the LCFs will continue to deliver HPC capabilities for large scale applications to ensure that the U.S. research community and DOE's industry partners continue to have access to the most capable supercomputing resources in the world. The OLCF will begin operating the new IBM Summit system to allow early science users to harness up to 200 petaflops of sustained performance while beginning preparation for an exaflop upgrade in 2021. The ALCF upgrade project will shift toward an advanced architecture, particularly well-suited for machine learning applications capable of more than an exaflop performance when delivered. This will impact site preparations and requires significant new non-recurring engineering efforts with the vendor to develop features that meet ECI requirements and that are architecturally diverse from the OLCF exascale system.

The NERSC will continue operations of the NERSC-8 supercomputer, named Cori, which has expanded the capacity of the facility to approximately 30 petaflops. To keep pace with the growing demand for capacity computing to meet mission needs, the FY 2018 Request supports site preparations and non-recurring engineering to deploy 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 accelerated timeline for the delivery of an exascale system, the Research and Evaluation Prototypes (REP) activity will continue to support the Computational Sciences Graduate Fellowship at $10,000,000 in FY 2018. 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, ASCR continues to support, within ASCR facilities funding, post-doctoral training program for high end computational science and engineering at the facilities through ASCR facilities funding. In addition, the three ASCR HPC user facilities will continue to prepare their users for future architectures.

ASCR also continues to support the future computing technologies testbed activity through REP. This research activity is focused on exploring the challenges and opportunities of quantum computing, a promising but currently experimental computing architecture. These efforts are in partnership with industry and the quantum research community.

In FY 2018, ESnet will continue to provide networking connectivity for large-scale scientific data flows. The last significant upgrade of the ESnet was in 2010 and some links are reaching the end of their life-span. In addition, the near-term delivery of exascale and sharply increased data rates from several other SC facilities means that the demand for data movement will exceed the cost effective capabilities of ESnet's rapidly aging technology. Therefore, the ESnet has an approved Mission Need Statement (Critical Decision-0) to initiate a significant upgrade project in FY 2018 that will incorporate new optical technologies and increase core capacity to more than one terabit (one trillion bits) per second - an increase of 2-10 times current capacity at significantly lower per-wave deployment costs. This activity will also support all of the critical research necessary to deploy these technologies with continued 99.999% reliability and enhanced cyber security protections.

Exascale Computing

The ASCR FY 2018 Budget Request includes $346,580,000 to significantly accelerate the development of exascale-capable computing systems and with the intention to deploy these systems in 2021 to meet national needs through the SC's Exascale Computing Project (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, weather and energy modeling, genomics and systems biology, as well as to support national security objectives and energy technology advances in DOE. Exascale systems' computational capabilities are also needed for increasing data-analytic and data-intense applications across the DOE science and engineering programs and other Federal organizations that rely on large-scale simulations, e.g., the Department of Defense and the National Institutes of Health. The importance of exascale computing to the DOE science programs is documented in individual requirements reviews for each SC program office. Because DOE partners with HPC vendors to accelerate and influence 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.

Exascale computing is a central component of long-term collaboration between the SC's ASCR program and the NNSA's Advanced Simulation and Computing Campaign (ASC) program to maximize the benefits of the Department's investments, avoid duplication and leverage the significant expertise across the DOE complex.

The primary goal of the ECI in SC is to develop the technologies needed to deploy an exascale computing capability to 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.

Recent information gathered from U. S. vendors and results from previous DOE investments in the Design Forward and Fast Forward pre-exascale activities have identified opportunities to accelerate our exascale efforts to keep pace with foreign competition. As shown in the following table, with the $346,580,000 SC ECI Request, SC intends to fund the acceleration of the research development of an exascale platform based on an advanced architecture, with the intention to deploy in 2021 at ANL, followed by a second exascale-capable system with a different advanced architecture at ORNL:

- $196,580,000 for the ECP project to accelerate research and the preparation of applications, develop a software stack for both exascale platforms, and support additional co-design centers in preparation for exascale system deployment in 2021.
- $150,000,000 in LCF activity to begin planning, non-recurring engineering, and site preparations for the intended deployment of at least one exascale system in 2021. Deployment of exascale systems will be through the LCFs as part of their usual upgrade processes.

This approach will reduce the risk of the project and expand the range of applications able to effectively utilize these capabilities in 2021.

FY 2018 Crosscuts ($K)

	ECI
Advanced Scientific Computing Research	346,580

Advanced Scientific Computing Research
Funding ($K)

	FY 2016 Enacted[a]	FY 2017 Annualized CR[b]	FY 2018 Request	FY 2018 vs FY 2016
Mathematical, Computational, and Computer Sciences Research				
Applied Mathematics	42,318	–	30,104	-12,214
Computer Science	39,160	–	29,296	-9,864
Computational Partnerships	34,336	–	41,268	+6,932
Next Generation Networking for Science	20,591	–	0	-20,591
SBIR/STTR	6,181	–	11,261	+5,080
Total, Mathematical, Computational, and Computer Sciences Research	**142,586**	**–**	**111,929**	**-30,657**
High Performance Computing and Network Facilities		–		
High Performance Production Computing	86,000	–	80,000	-6,000
Leadership Computing Facilities	182,517	–	249,321	+66,804
Exascale (non-add)	*0*	*0*	*(150,000)*	*+150,000*
Research and Evaluation Prototypes	156,820	–	24,452	-132,368
High Performance Network Facilities and Testbeds	38,040	–	45,000	+6,960
SBIR/STTR	15,037	–	14,728	-309
Total, High Performance Computing and Network Facilities	**478,414**	**–**	**413,501**	**-64,913**
Exascale Computing				
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	0	0	196,580	+196,580
Total, Advanced Scientific Computing Research	**621,000**	**619,819**	**722,010**	**+101,010**

SBIR/STTR funding:

- FY 2016 Enacted: SBIR $18,450,000 and STTR $2,768,000
- FY 2018 Request: SBIR $22,785,000 and STTR $3,204,000

[a] The FY 2016 Enacted level includes SBIR and STTR and reflects updates through the end of the fiscal year.
[b] FY 2017 Annualized CR amounts reflect the P.L. 114-254 continuing resolution level annualized to a full year. These amounts are shown only at the congressional control level and above, below that level, a dash (-) is shown.

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

	FY 2018 vs FY 2016
Mathematical, Computational, and Computer Sciences Research: The Computer Science and Applied Mathematics activities will increase their emphasis on data-intensive science and initiate new efforts in machine learning to increase the impact of data generated in extreme-scale simulations and by Office of Science user facilities. The Computational Partnerships activity will continue to infuse the latest developments in applied math and computer science into the strategic applications of the Office of Science to get the most out of the leadership computing systems. Investments will increase for future computing technologies such as quantum information systems in partnership with BES, BER, and HEP. The Next Generation Networking for Science activity will be eliminated. Collaboratory efforts currently supported by this activity will be supported by computational partnerships to strengthen the interconnectivity of these efforts. Networking-related R&D will be supported within the High Performance Network Facilities and Testbeds and will focus on the unique needs of the ESnet in support of a significant upgrade to handle the increased data flows from the exascale systems and other Office of Science user facilities.	**-30,657**
High Performance Computing and Network Facilities: Increased facilities funding initiates activities with the intention to deploy an exascale system at the ALCF in 2021 and to begin preparations at the OLCF for an exascale system that is architecturally distinct to follow in 2022. Funding also supports operations, including increased power costs at all facilities. Research and Evaluation Prototypes funding will support quantum computing testbeds and the Computational Sciences Graduate Fellowship at current levels. High Performance Network Facilities and Testbeds supports initiation of a significant upgrade, ESnet 6, to ensure that networking and data transfer capacity keep pace with increasing demand as new facilities come online, including the first exascale machine in 2021.	**-64,913**
Exascale Computing: Delivery of at least one exascale system based on an advanced architecture is planned for 2021, with a second exascale system with a different architecture to also follow in 2022. The FY 2018 budget also supports accelerated application and software development for both architectures, together with additional co-design centers, applications, partnerships with the vendors, and testbeds to ensure development of hardware and software keep pace.	**+196,580**
Total, Advanced Scientific Computing Research	**+101,010**

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 HPC 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. 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 2018, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Early Exascale Investments Pay Off for U.S. Computing Industry. An exascale architecture with enhanced integration of processors and memory developed as part of the FastForward 2 program has already improved both performance and power efficiency in one vendor's data centers and graphics products. The FastForward funding allowed the vendor to take a holistic approach that includes detailed analysis of innovative computing solutions, acceleration of key hardware and software technologies, and co-design with DOE laboratories that would not otherwise have been possible. Additional anticipated benefits include development of power management techniques that will allow processors to optimize power and performance based on application and user requirements; better programming environments, tools, and libraries to simplify parallel programming and improve performance for large-scale heterogeneous systems; and outstanding opportunities for students and early-career scientists to work in industry in advanced technology areas. The partnerships between U.S. vendors and DOE simultaneously advances DOE's science and national security missions and U.S. competitiveness by helping vendors develop computing infrastructure solutions that enhance the business case for deploying platforms for HPC, data analytics, and other high-growth markets. Studies have shown that growth in computers and microelectronics yields the biggest returns in both GDP and high quality jobs[a].

Reducing the Damage Caused by Industrial Fires. Warehouse fires are the leading cause of commercial property damage, responsible for 40% of all industry property loss at a cost of approximately $188 million per year. Understanding how fires spread has the potential to save both business owners and insurance companies hundreds of millions of dollars. However, some of the most destructive fires – those that take place in mega-warehouses with ceilings up to 100 ft. high and a footprint in excess of 100,000 sq. ft.– are among the most difficult to study because they cannot be replicated in a test facility. To solve this problem, one of the world's largest commercial and industrial insurance companies partnered with the Oak Ridge Leadership Computing Facility to adapt an open-source fluid dynamics code to include the complex processes that occur during an industrial fire, including soot formation and sprinkler spray dynamics. After running their high-resolution FireFOAM code on the Oak Ridge Leadership Computing Facility's Titan machine to learn how to stack storage boxes on pallets to impede the spread of horizontal flames, the team incorporated the SciDAC-developed Adaptable I/O System (ADIOS) into FireFOAM to improve its efficiency in moving data on and off the supercomputer. The new and improved code is now being used to simulate other commodities stored in warehouses, starting with large paper rolls. Both the results and the code are shared publicly to promote the improvement of fire protection standards across industry.

New Algorithm for Incompressible Fluid Flow Modeling. The Navier-Stokes equations for fluid flow are used for applications ranging from special effects in movies to industrial design. By solving these complex equations, researchers can gain insight into how fast a fluid is moving through its environment, how much pressure it is under, and the forces it exerts on its surroundings. Using the NERSC, a researcher at LBNL reformulated the incompressible Navier-Stokes equations to make them more amenable to numerical computation and overcome challenges in resolving the intricate fluid dynamics near moving boundaries such as bubbles, swimming organisms, and surface waves. The new algorithms capture changing small-scale features near these boundaries with unprecedented detail and provide additional information about how the tiny features influence the fluid far away from the boundary. The resulting cheaper computational models and increased resolution capabilities will allow researchers to study even more complex phenomena such as blood flow in the pumping

[a] A Nov 2016 Deutsche Bank Report (Zhang and Zeng) estimated that a 10% reduction in the U.S. trade deficit in "computers and electronics" would add $18 billion to the U.S. economy.

heart, the ejection of ink droplets in consumer inkjet printers, and optimal propeller design. This work is described in the June 10, 2016 issue of Science Advances[a].

Modeling Mini-Proteins to Design More Effective Drugs with Minimal Side Effects. Therapeutic drugs typically work by attaching themselves to, and thus disabling, disease-inducing molecules in the human body. The perfect drug has a stable structure that will move easily through the body to bind to its intended target and nothing else. Many of the therapeutic drugs in common use today are small molecules that dissolve readily in the body but bind to unintended targets in a manner that can lead to harmful side effects. Drugs based on protein molecules that have a highly specific structure can minimize side effects at the cost of being too large to pass through membranes that are their entry route into the human body. Peptides – molecules derived from the same building blocks as proteins but many times smaller – may strike the perfect balance, as long as they include an additional man-made component to stabilize their structure. To evaluate a peptide for use as a therapeutic drug, researchers must perform hundreds of millions of simulations to explore how the peptide will fold and bind to disease-inducing molecules. Researchers at the University of Washington, funded by NIH and NSF, have put the ALCF's more than 780,000 cores to use to refine their technique for evaluating and optimizing over thousands of different peptide designs. The team's long-term goal is to create a database of peptides to assist in future drug design efforts. This work was published in Nature Magazine in October 2016[b].

Network Pilot Quintuples Data Transfer Rates. Over half of the researchers using the General Medical Sciences and National Cancer Institute Structural Biology (GM/CA) Facility at ANL's Advanced Photon Source access the facility from their home institutions. While remote access allows a wider range of researchers to use the facility's intense, tunable x-ray microbeams, remote users often cannot analyze their data until after their beamtime is over because it takes too long to move the large data files to a computer capable of processing them. A pilot program to implement the Science DMZ – an ESnet network model tailored to the needs of high-performance science applications – at the GM/CA facility demonstrated data transfer rates over five times faster between the facility and Purdue University for most file sizes. Additional tests are planned at five more remote sites. Integrating the lessons learned in this ESnet pilot into regular operations will simplify the facility's support for the beamline and lead to increased scientific productivity.

[a] *Interfacial gauge methods for incompressible fluid dynamics*. Science Advances 10 Jun 2016: Vol. 2, no. 6, e1501869 DOI: 10.1126/sciadv.1501869

[b] http://www.nature.com/nature/journal/v538/n7625/abs/nature19791.html

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 workflows to efficiently and effectively harness the potential of today's HPC systems and advanced networks for science and engineering applications;
- operating systems, data management, analyses, machine learning, 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; 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, data movement, resilience, and machine learning in digital computing, and exploratory research in future computing paradigms that have the potential to revolutionize scientific computing in the post-exascale era.

Deriving scientific insights from the vast amounts of data flowing from SC user facilities as well as the output of extreme-scale simulations will require a sophisticated tool suite for data manipulation, visualization, pattern recognition, and analysis. Data-intensive science additionally requires a focused research effort to develop the necessary theories, software tools, and technologies to manage the full data lifecycle from generation or collection through capturing the historic record of the data, archiving, and sharing them. In FY 2018, ASCR's research program will increase its emphasis on data-intensive science, including new efforts to exploit machine learning and other adaptive algorithms to enhance the impact of scientific data generated across SC.

<u>Applied Mathematics</u>
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 complex physical and engineered 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 future HPC systems. High-fidelity modeling and simulation require a number of new algorithmic techniques and strategies supported by this activity, including adaptive algorithms and machine learning, 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.

<u>Computer Science</u>
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 for 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 data management and analysis.

Significant innovation in computer science is needed to realize the computational and data-analytic potential of future HPC systems and other scientific user facilities in a timeframe consistent with their anticipated availability. There will be an increased emphasis on data-intensive science challenges with particular attention to adaptive algorithms and machine learning, the intersection with exascale computing challenges, and the unique needs of DOE scientific user facilities including data management and cyber security. 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. 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 HPC, 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 and many-core architectures and machine learning techniques.

SciDAC 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. The SciDAC institutes bridge core research efforts in algorithms, methods, software, and tools with the need of the SciDAC applications supported in partnership with the other SC programs. Current SciDAC applications include chemistry, materials science, fusion research, high energy physics, nuclear physics, astrophysics, earth systems modeling, and accelerator physics. In FY 2018 these efforts also include the collaboratory partnerships previously supported by the Next Generation of Networking for Science. These efforts enable large distributed research teams to share data and develop tools for real-time analysis of the massive data flows from Office of Science scientific user facilities.

In addition to SciDAC, the Computational Partnerships activity supports interdisciplinary teams in partnership with BES, BER, and HEP to develop algorithms and applications targeted for future computing platforms, including quantum information systems.

Activities and Explanation of Changes

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
Mathematical, Computational, and Computer Sciences Research $142,586,000	**$111,929,000**	**-$30,657,000**
Applied Mathematics $42,318,000	**$30,104,000**	**-$12,214,000**
Applied Mathematics continued efforts to develop new algorithmic techniques and strategies that extract scientific advances and engineering insights from massive data for DOE missions. Applied Mathematics addressed many of the challenges of exascale including: advanced solvers, uncertainty quantification, algorithmic resilience, and strategies for reducing global communications.	Applied Mathematics will continue its core programs in new algorithmic techniques and strategies that extract scientific advances and engineering insights from massive data for DOE missions. Adaptive algorithms and machine learning will be added to the suite of tools under development for optimizing the scientific output of data-intensive programs across SC.	Decrease reflects transfer of exascale related funding to the SC-ECP and does not constitute a change in scope for these activities.
Computer Science $39,160,000	**$29,296,000**	**-$9,864,000**
Computer Science continued efforts to develop software, new programming models and metrics for evaluating system status. This activity primarily focused on addressing the challenges of exascale and data-intensive science and emphasized efforts to promote ease of use, increased parallelism, energy efficiency, and reliability, and ensured 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. Activities to support development of future computing technologies will also continue, and a new effort will be initiated to exploit machine learning techniques to better understand data generated both by HPC simulations and SC facilities.	Decrease reflects transfer of exascale related funding to the SC-ECP and does not constitute a change in scope for these activities.

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
Computational Partnerships $34,336,000	**$41,268,000**	**+$6,932,000**
The SciDAC Institutes were 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 LCFs, the NERSC, and similar world-class computing facilities over the next five years.	The SciDAC institutes continue to play a key role in assisting DOE mission critical applications to effectively use ASCR's existing production and LCFs while the newly-awarded fourth-generation SciDAC partnerships focus on preparing applications to harness the potential of ASCR's planned upgrades to its computing facilities. Partnerships initiated in FY 2017 to explore potential impacts of future computing technologies across SC continue. In FY 2018, this effort also supports collaboratory partnerships previously included in Next Generation Networking for Science.	Increase reflects the transfer of some funding from the Next Generation Networking for Science activity to support collaboratory partnerships. Funding also increases to support interdisciplinary partnerships to develop algorithms and applications for quantum information systems.
Next Generation Networking for Science $20,591,000	**$0**	**-$20,591,000**
The Next Generation Networking for Science activity continued 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 allowed 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 be eliminated through consolidation with other ASCR activities and successful completion of several projects in FY 2018.	Collaboratory efforts, currently supported by this activity, will be supported by the computational partnerships activity to strengthen the interconnectivity of these efforts. Networking R&D will be supported by the High Performance Network Facilities and Testbeds activity and will focus on the unique needs of the ESnet in support of a significant upgrade of the ESnet to address the increased data flows from the exascale systems and other Office of Science user facilities.
SBIR/STTR $6,181,000	**$11,261,000**	**+$5,080,000**
In FY 2016, SBIR/STTR funding is set at 3.45% of non-capital funding.	In FY 2018, SBIR/STTR funding is set at 3.65% of non-capital funding.	

Description

The High Performance Computing and Network Facilities subprogram delivers forefront computational and networking capabilities. These include high performance production computing at NERSC at LBNL and LCFs at ORNL and ANL. These computers, and the other SC 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 ESnet. Finally, 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 (REP) activity addresses the challenges of next generation computing systems. By actively partnering with the research community, including industry, on the development of technologies that enable next-generation machines, ASCR ensures that commercially available architectures serve the needs of the scientific community. The REP activity also prepares researchers to effectively use future generations of scientific computers, including novel technologies, and seeks to reduce risk for future major procurements.

Allocation of computer time at ASCR facilities follows the peer-reviewed and 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, earth systems modeling, materials, high energy and nuclear physics, fusion, and biology. NERSC users come from nearly every state in the U.S., with about 49% based in universities, 46% in DOE laboratories, and 5% in other government laboratories and industry. NERSC's large and diverse user base requires an agile support staff to aid users entering the HPC arena for the first time, as well as those preparing codes to run on the largest machines available at NERSC and the LCFs. In FY 2015, NERSC moved into the new Computational Research and Theory building located on the LBNL campus.

NERSC is a vital resource for the SC 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 SC 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 was one of the most powerful computers in the world for scientific research and was ranked number three on the November 2016 Top 500 list, just below the most powerful supercomputers in China.[a] The FY 2018 upgrade of this facility to a 200 pf system will challenge the leadership of

[a] http://www.top500.org/lists/2016/11/

the world's fastest systems. Early science applications at the OLCF, including large eddy simulation of turbulent combustion in complex geometries, quantum Monte Carlo simulations for the study and prediction of materials properties, heavy element chemistry, models of astrophysical explosions, dynamical simulations of magnetic fields in high-energy-density plasmas, molecular design of next-generation nanochemistry for atomically precise manufacturing, simulation of cellular and neural signaling, simulations of neutron transport in fast-fission reactor cores, and earthquake simulations, are scaling to make effective use of the new capability. 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. These efforts often result in upgrades to in-house computing resources at these U.S. companies.

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 REP activity. This HPC system achieves high performance with relatively lower electrical power consumption than other current petascale computers. The ALCF also operates an 8.5 pf Intel-based machine (Theta) to prepare their users for the ALCF-3 upgrade in 2019-2020.

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 of aircraft engines.

The demand for 2016 INCITE allocations at the LCFs outpaced the available resources by a factor of two with growth expected to sharply increase upon the availability of upgrades.

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 future computing technologies "Beyond Moore's Law", specifically in the area of quantum computing.

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

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. ESnet underpins 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. ESnet supports the data requirements of all SC facilities, including the increased bandwidth and high-traffic links added in FY 2017 as well as transatlantic access to Large Hadron Collider data and direct science engagement efforts to improve end-to-end network performance between DOE facilities and U.S. universities. 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. The Request includes additional funding to build a next-generation network, ESnet 6, which will meet the growing data needs of the SC facilities, including the intention to deploy the first exascale machine in 2021. This activity will also support all of the critical research necessary to deploy these technologies with continued reliability with enhanced cyber security protections.

Activities and Explanation of Changes

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
High Performance Computing and Network Facilities $478,414,000	**$413,501,000**	**-$64,913,000**
High Performance Production Computing $86,000,000	**$80,000,000**	**-$6,000,000**
Supported 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.	The NERSC-8 system, named “Cori” after Nobel Laureate Gerty Cori, will continue production operations. Demand for production computing for the SC programs continues to grow along with system capability and the rapid increase in data from experiments. In FY 2018, preparation for the 2020 delivery of NERSC-9, which will provide three to five times the capacity of NERSC-8, will continue.	Decrease reflects completion of site preparations for NERSC-8. Operation of NERSC-8 and the NERSC-9 acquisition continue as planned.
Leadership Computing Facilities $182,517,000	**$249,321,000**	**+$66,804,000**
Supported 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 included lease payments, power, and user support. Also supported 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.	Operation will continue at both LCF facilities while upgrades will proceed as planned. The OLCF will install, test, provide early science access, and transition the new IBM hybrid supercomputer, called Summit, to operations in early FY 2018. This upgrade will provide 200 pf of computing capability, or approximately five times the capability of the previous system, Titan. OLCF will also begin site preparations to enable deployment of an exascale system as early as 2021. In FY 2018, the ALCF will continue to provide access to the 8.5-pf Intel Xeon interim system, called Theta, deployed in early FY 2017 to transition ALCF users to the new architecture. The ALCF will begin site preparations and significant nonrecurring engineering efforts to deploy a novel architecture capable of delivering more than an exaflop of computing capability in 2022.	Increase provides for planning and the initiation of site preparation activities and non-recurring engineering to ready both facilities to deploy exascale systems, with distinct architectures in 2021-2. This approach reduces the risk of the exascale initiative and broadens the community able to of utilize these new capabilities.
Leadership Computing Facility at ANL: $77,200,000	$100,000,000	+$22,800,000

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
Leadership Computing Facility at ORNL: $105,317,000	$149,321,000	+$44,004,000
Research and Evaluation Prototypes $156,820,000	**$24,452,000**	**-$132,368,000**
REP supported efforts to improve the energy efficiency and reliability of critical technologies such as memory, processors, network interfaces and interconnects. REP competitively selected R&D partnerships with U.S. vendors to initiate the design and development of compute node and system designs suitable for exascale systems, which was an essential component of the Department's exascale computing plan and a key step in the vendor's productization efforts. To emphasize the vital importance of the CSGF program to the ASCR facilities and to our exascale goals, Research and Evaluation Prototypes supported the program at $10,000,000.	Availability of experienced and knowledgeable workforce issues continues to be of vital importance to ASCR's current and planned facilities. The CSGF program will play an increasingly important role as the Exascale Initiative progresses and future computing technologies mature. Therefore, support for the CSGF within REP continues at $10,000,000 in FY 2018. This activity will also provide increased support for the future computing technology testbed focused on quantum computing established in FY 2017.	Decrease reflects transfer of exascale related funding to the SC-ECP and does not constitute a change in scope for these activities while there is an increase to support additional quantum computing testbed activities.
High Performance Network Facilities and Testbeds $38,040,000	**$45,000,000**	**+$6,960,000**
ESnet operated 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.	The Request supports operations and maintenance of the network and continued development of 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 networking R&D, previously supported by the Next Generation Networking for Science (NGNS) activity, 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. ESnet was last upgraded in 2010 and some technology is no longer supported by the vendor. Additional funds are therefore requested for the upgrade to ESnet 6, which will provide a network for scientific	Increase supports the ESnet 6 upgrade required to meet SC data requirements through the mid-2020s and supports related networking research previously supported by the NGNS activity.

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
	data transfer with the capacity, reliability and resilience, and flexibility to meet the needs of the Office of Science facilities and research community through the mid-2020s.	
SBIR/STTR $15,037,000	**$14,728,000**	**-$309,000**
In FY 2016, SBIR/STTR funding is set at 3.45% of non-capital funding.	In FY 2018, SBIR/STTR funding is set at 3.65% of non-capital funding.	

Description

The Office of Science Exascale Computing Project (SC-ECP) in the Exascale Computing subprogram captures the research aspects of ASCR's participation in the U. S. Department of Energy's Exascale Computing Initiative (ECI), to ensure the hardware and software R&D, including applications software, for a exascale-capable system is completed in time to meet the scientific and national security mission needs of the DOE in 2021. The deployment of these systems, including necessary site preparations and non-recurring engineering, is supported by the Leadership Computing Facilities activity that will ultimately house and operate the exascale systems. The ECI will execute a program, joint between SC and NNSA, to develop and deploy an exascale-capable computing system with an emphasis on sustained performance for relevant applications and analytic computing to support DOE missions.

The SC-ECP supports R&D for the development 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*, tailored for this fast-paced research effort and similar to that which has been used by SC for the planning, design, and construction of all of its major computing projects, including the LCFs at Argonne and Oak Ridge National Laboratories and NERSC at LBNL.

The FY 2018 Request includes $196,580,000 for SC-ECP. These funds support the preparation of applications; and the development of a software stack for both platforms. Funding also supports additional co-design centers, vendor partnerships, and testbeds in preparation for the intended deployment of an exascale system in 2021. Deployment of exascale systems will be through the LCFs as part of their usual upgrade processes. $150,000,000 of additional ECI funding is provided in the LCF activity to begin planning, non-recurring engineering, and site preparations for the intended delivery of at least one exascale system as early as 2021.

Activities and Explanation of Changes

FY 2016 Enacted	FY 2018 Request	Explanation of Changes FY 2018 vs FY 2016
17-SC-20 Office of Science Exascale Computing Project (SC-ECP) $0	**$196,580,000**[a]	**+$196,580,000**[a]
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 2018 funding will accelerate application and software stack development in preparation for the intended delivery of an exascale system in 2021.	The increase reflects revision of ECP for the intended delivery of at least one exascale system as early as 2021. Application and software development is accelerated, and additional codesign centers have been added to the project.

[a] In addition, $150,000,000 of ECI funding is requested within the Leadership Computing Facilities activity to begin planning, non-recurring engineering, and site preparations for intended deployment of at least one exascale system in 2021.

Office of Science
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.

	FY 2016	FY 2017	FY 2018
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	Fund two teams to develop exascale node designs.	Identify at least one multi-institutional team to develop new mathematics for DOE mission focused grand challenges at the nexus of multiple computational sub-domains such as data-driven	Support at least two machines learning efforts in both applied mathematics and computer science.

		discovery, multiscale modeling, uncertainty quantification, and adaptive algorithms.	
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 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Request	FY 2018 vs FY 2016
Capital operating expenses						
Capital equipment	n/a	n/a	5,700	–	10,000	+4,300

Funding Summary ($K)

	FY 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Request	FY 2018 vs FY 2016
Research	293,225	–	321,700	+28,475
Scientific user facility operations	306,557	–	374,321	+67,764
Other	21,218	–	25,989	+4,771
Total, Advanced Scientific Computing Research	**621,000**	**619,819**	**722,010**	**+101,010**

[a] FY 2017 Annualized CR amounts reflect the P.L. 114-254 continuing resolution level annualized to a full year. These amounts are shown only at the congressional control level and above, below that level, a dash (-) is shown.

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 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Request	FY 2018 vs FY 2016
TYPE A FACILITIES				
NERSC	**$86,000**	**$–**	**$80,000**	**-$6,000**
Number of Users	5,608	–	6,000	+392
Achieved operating hours	N/A	–	N/A	
Planned operating hours	8,585	–	8,585	+0
Optimal hours	8,585	–	8,585	+0
Percent optimal hours	N/A	–	N/A	
Unscheduled downtime hours	N/A	–	N/A	

[a] FY 2017 Annualized CR amounts reflect the P.L. 114-254 continuing resolution level annualized to a full year. These amounts are shown only at the congressional control level and above, below that level, a dash (-) is shown.

	FY 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Request	FY 2018 vs FY 2016
OLCF	**$105,317**	–	**$149,321**	**+$44,004**
Number of Users	1,064	–	1,064	+0
Achieved operating hours	N/A	–	N/A	
Planned operating hours	7,008	–	7,008	+0
Optimal hours	7,008	–	7,008	+0
Percent optimal hours	N/A	–	N/A	
Unscheduled downtime hours	N/A	–	N/A	
ALCF	**$77,200**	–	**$100,000**	**+$22,800**
Number of Users	1,434	–	1,434	+0
Achieved operating hours	N/A	–	N/A	
Planned operating hours	7,008	–	7,008	+0
Optimal hours	7,008	–	7,008	+0
Percent optimal hours	N/A	–	N/A	
Unscheduled downtime hours	N/A	–	N/A	
ESnet	**$38,040**	–	**$45,000**	**+$6,960**
Number of users[b]	N/A	–	N/A	
Achieved operating hours	N/A	–	N/A	
Planned operating hours	8,760	–	8,760	+0
Optimal hours	8,760	–	8,760	+0
Percent optimal hours	N/A	–	N/A	
Unscheduled downtime hours	N/A	–	N/A	

[a] FY 2017 Annualized CR amounts reflect the P.L. 114-254 continuing resolution level annualized to a full year. These amounts are shown only at the congressional control level and above, below that level, a dash (-) is shown.

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

	FY 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Request	FY 2018 vs FY 2016
Total Facilities	**$306,557**	–	**$374,321**	**+$67,764**
Number of Users[b]	8,106	–	8,498	+392
Achieved operating hours	N/A	–	N/A	
Planned operating hours	31,361	–	31,361	+0
Optimal hours	31,361	–	31,361	+0
Percent of optimal hours[c]	N/A	–	N/A	
Unscheduled downtime hours	N/A	–	N/A	

Scientific Employment

	FY 2016 Enacted	FY 2017 Annualized CR[a]	FY 2018 Estimate	FY 2018 vs FY 2016
Number of permanent Ph.D.'s (FTEs)	584	–	601	+17
Number of postdoctoral associates (FTEs)	146	–	200	+54
Number of graduate students (FTEs)	460	–	486	+26
Other scientific employment (FTEs)[d]	247	–	257	+10

[a] FY 2017 Annualized CR amounts reflect the P.L. 114-254 continuing resolution level annualized to a full year. These amounts are shown only at the congressional control level and above, below that level, a dash (-) is shown.

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

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

[d] 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

This Project Data Sheet (PDS) is an update of the FY 2017 PDS and does not include a new start.

The FY 2018 Request for SC-ECP is $196,580,000. The most recent DOE O 413.3B approved Critical Decision (CD) is CD-1/3A, Approve Alternative Selection and Cost Range and Approve phase one funding of hardware and software research projects and application development, was approved on January 3, 2017. The estimated Total Project Cost (TPC) range of the SC-ECP is $1.0 billion to $2.7 billion.

Summary

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 a partnership between the Office of Science (SC) and the National Nuclear Security Administration (NNSA) and addresses Department of Energy's (DOE) science and national security mission requirements.

In FY 2017, the SC 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 research and development (R&D) activities required for the development of exascale-capable computers. Other activities related to the ECI but outside of the scope of the R&D activities leading to exascale-capable computers are not within the SC-ECP, though they do remain in the scope of the ECI. In FY 2018 these include $150,000,000 to support the initiation of planning, site preparations, and non-recurring engineering at the Leadership Computing Facilities (LCFs) where the exascale machines will be housed and operated. With all of these activities and funding, DOE intends to accelerate delivery of at least one exascale-capable system in 2021. Supporting parallel development at both LCFs will reduce the overall risk of the project and broaden the range of applications able to utilize this new capability. This PDS is for the SC-ECP only; prior-year activities related to the SC-ECP are also included.

In FY 2018, SC-ECP funding will support project management; development of project documentation; conduct of co-design activities with a representative subset of mission applications; R&D of exascale systems software and tools needed for exascale programming; and vendor partnerships.

2. Critical Milestone History and Schedule

(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
FY 2018	07/28/2016	TBD	01/03/2017	4Q FY 2019	3Q FY 2019	4Q FY 2019	N/A	4Q FY 2023

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 $1.0 billion and $2.7 billion. The cost range will be updated and a project baseline (scope, schedule, and cost) will be established at CD-2.

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-capable 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-capable 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 three focus areas:

- *Hardware Technology:* The Hardware Technology focus area supports vendor-based research and development activities required to deploy at least two exascale-capable 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-capable 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.

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 SC 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 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 three years of the project, project activities will focus on hardening the application and the system stack software, and additional hardware technologies investments 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. Financial Schedule

	(dollars in thousands)		
	Appropriations	Obligations	Costs
Total Estimated Cost (TEC) (Hardening of Applications Development System Software Technology, Hardware Technology)			
FY 2016[a]	0	0	0
FY 2019– FY 2023	390,000	390,000	390,000
Total, TEC	390,000	390,000	390,000
Other project costs (OPC) (Research for Application Development , System Software Technology and Hardware Technology)			
FY 2016	157,944	157,944	12,500
FY 2017	164,000	164,000	228,000
FY 2018	196,580	196,580	350,444
FY 2019 – FY 2023	245,000	245,000	172,580
Total, OPC	763,524	763,524	763,524
Total Project Costs (TPC)			
FY 2016	157,944	157,944	12,500
FY 2017	164,000	164,000	228,000
FY 2018	196,580	196,580	350,444
FY 2019 – FY 2023	635,000	635,000	562,580
Total, TPC	1,153,524	1,153,524	1,153,524

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

6. Project Cost Estimate

The SC-ECP will be baselined at CD-2. 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)			
Application Development	225,000	TBD	N/A
Production Ready Software	86,000	TBD	N/A
Hardware Partnerships	79,000	TBD	N/A
Total, TEC	390,000	TBD	N/A
Other Project Costs (OPC) (Research)			
Planning/Project Mgmt	118,000	8,000	N/A
Application Development	269,630	85,000	N/A
Software Research	121,423	87,000	N/A
Hardware Research	254,471	131,894	N/A
Total OPC	763,524	TBD	N/A
Total, TPC	1,153,524	TBD	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
FY 2018	TEC	0	0	0	0	175,000	145,000	70,000	390,000
	OPC	157,944	164,000	196,580	189,000	14,000	14,000	28,000	763,524
	TPC	157,944	164,000	196,580	189,000	189,000	159,000	98,000	1,153,524

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

8. Related Operations and Maintenance Funding Requirements

System procurement activities for the exascale-capable computers are not part of the SC-ECP. The exascale-capable computers will become part of existing facilities and operations and maintenance funds and will be included in the ASCR facilities' operations budget. In the FY 2018 President's Request, $150,000,000 is included in the Argonne Leadership Computing Facility and the Oak Ridge Leadership Computing facilities budgets to begin planning non-recurring engineering and site preparations for the delivery and deployment for the exascale systems. These funds are included in ECI but not SC-ECP.

Start of Operation	2022
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.

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 three years duration will support finalizing applications and system software.