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, development, and deployment of new computing systems along with the applied mathematics and software technologies to effectively use the leading edge systems. The benefits of U.S. computational leadership have been enormous – huge gains in increasing workforce productivity, accelerated 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, which 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] U.S. dominance in computing is under threat from significant investments in Asia and Europe. This is happening at a time when advances in computing capabilities are becoming increasingly costly and risky. We cannot afford to fall further behind in an area that is critical to the American prosperity because computing impacts our national security, every sector of our economy, and every field of science and engineering. Therefore, this FY 2019 Request increases our investments in the Department of Energy (DOE) Exascale Computing Initiative (ECI) to lower the risk of the project and deliver at least one exascale-capable system in 2021 and a second system in the 2021-2022 timeframe. In addition, the FY 2019 Request includes support for long-term investments in future computing technologies such as quantum information systems and artificial intelligence, areas where Asia and Europe are also investing heavily.

The 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 at the end of the current technology.

Maximizing the benefits of U.S. leadership in 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 in the National Strategic Computing Initiative (NSCI), DOE will sustain and enhance its support for HPC research, development, and deployment as part of a coordinated Federal strategy guided by five NSCI objectives:

- Accelerate delivery of a capable exascale computing system that 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.
- Increase coherence between the technology base used for modeling and simulation and the one used for data analytic computing.
- Establish, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post- Moore's Law era").

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

- 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.
- Develop an enduring public-private collaboration to ensure that the benefits of the research and development (R&D) advances are maximized and to the greatest extent, shared between the United States Government and industrial and academic sectors.

Within the context of coordinated NSCI strategy, the DOE Office of Science (SC) initiated a long term research program in quantum information science to prepare for the "post-Moore's Law era." In addition, SC and the DOE National Nuclear Security Administration (NNSA) are partnering on the ECI to overcome key exascale challenges in parallelism, energy efficiency, and reliability, leading to deployment of a diverse set of exascale systems in the 2021-2022 timeframe. The ECI's goal for an exascale-capable system is a fifty-fold increase in sustained performance over today's Titan system at Oak Ridge National Laboratory (ORNL), with applications that address next-generation science, engineering, and data problems. The ECI focuses on delivering advanced simulation through an exascale-capable computing program, which emphasizes sustained performance on science and national security mission applications and increased convergence between exascale and large-data analytic computing.

The SC FY 2019 Request funds two components of the ECI: planning, site preparations, and non-recurring engineering (NRE) at the Leadership Computing Facilities (LCF) to prepare for deployment of at least one exascale system in 2021, and the ASCR-supported Office of Science Exascale Computing Project (SC-ECP), first proposed in the FY 2017 Request, which includes only the R&D activities required to develop exascale-capable computers.

The scope of the SC-ECP has three areas which are focused on increasing the convergence of big compute and big data and creating a holistic high performance computing ecosystem:

- *Hardware Technology and Integration:* The goal of the Hardware Technology and Integration focus area is to integrate the delivery of SC-ECP products on targeted systems at leading DOE computing facilities.
- *System Software Technology:* The goal of the System Software Technology focus area is to produce a vertically integrated software stack to achieve the full potential of exascale computing including the software infrastructure to support large data management and data science for the DOE at exascale; and
- *Application Development:* The goal of the Application Development focus area is to develop and enhance the predictive capability of applications critical to DOE that involves: working with scientific and data-intensive grand challenge application areas to address the challenges of extreme parallelism, reliability and resiliency, deep hierarchies of hardware processors and memory, scaling to larger systems.

The SC-ECP is managed following the principles of DOE Order 413.3B, *Program and Project Management for the Acquisition of Capital Assets*, which SC uses 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 is conducted via a Project Office established at ORNL because of its considerable expertise in developing computational science and engineering applications; in managing HPC facilities, both for the Department and for other federal agencies; and experience in managing distributed, large-scale scientific research projects, such as the Spallation Neutron Source project. A Memorandum of Agreement is in place 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 2019 Request

The FY 2019 Request for ASCR increases our investments in the ECI to lower project risk to deliver at least one exascale-capable system in 2021. To ensure ASCR is meeting the HPC mission needs of the Office of Science during and after the exascale project, this Request prioritizes basic research for data intensive science, including machine learning and future computing technologies, and increases support for ASCR's Computational Partnerships with a focus on developing new strategic partnerships in quantum computing and data intensive applications. Funding for the ASCR facilities supports continued operations of the 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 will also support site preparations and NRE activities at the LCFs in support of a 2021 delivery of at least one exascale computing system.

<u>Mathematical, Computational, and Computer Sciences Research</u>
Recognizing that Moore's Law, where 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 science, that are not based on silicon microelectronics. In FY 2019, ASCR will continue the research and computational partnerships with the other SC program offices aimed at understanding the challenges that quantum information science 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 2019, 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. In addition, fundamental changes to experimental and observational sciences are presenting new data and computing challenges along with exciting opportunities to deliver transformative capabilities to researchers. These changes are brought about by a confluence of factors that include, for example, the tremendous increases in data volumes and streaming data rates from experimental and observational devices; opportunities to combine data sources for new insights; a need for real-time feedback or automated steering of experiments; and a need to find new ways to derive insights from increasingly complex, heterogeneous information resources. Potential solutions to address these challenges rely on providing researchers with efficient and timely access to data and computing resources; effective use of new computing hardware; and new tools for analysis of big data. The FY 2019 Request supports fundamental research in computer science and applied mathematics to address these challenges, as well as partnerships among computer scientists, applied mathematicians, networking specialists, and domain scientists to address end-to-end analysis needs. The Request places a particular emphasis on machine learning for discovery and decision support.

The Scientific Discovery through Advanced Computing (SciDAC) computational partnerships, which were recompeted in FY 2017, use the software, tools, and methods developed by these core research efforts. 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. In addition, new strategic partnerships in quantum computing and data intensive applications will advance capabilities and broaden the impact of these efforts. SciDAC investments are 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 was combined with other activities in FY 2018. Collaboratory efforts previously supported by this activity are now supported by computational partnerships to strengthen the interconnectivity of these efforts. Other networking-related R&D is now supported within the Computer Science activity. In FY 2019, these efforts will include research in new efforts in quantum networking.

<u>High Performance Computing and Network Facilities</u>
In FY 2019, 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 ALCF upgrade project will continue to focus on deployment of a system capable of delivering more than an exaflop performance in 2021. Site preparations continue along with NRE efforts with the vendor to develop features that meet ECI requirements. The OLCF will operate the IBM Summit system to allow users to harness up to 200 petaflops of sustained performance while beginning preparation for an exascale upgrade in the 2021-2022 timeframe that is architecturally diverse from the ALCF 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 address growing demand for capacity computing to meet mission needs, the FY 2019 Request supports activities for an upgrade to NERSC-9, which will have three to five times the capacity of NERSC-8 in

2020. The Request also supports site preparation activities for the NERSC-9 upgrade, such as increased power and cooling capacity, and NRE to ensure the new computing system meets the need of the diverse NERSC user community.

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. 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 the 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 will also continue 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 2019, 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) for the ESnet-6 upgrade in FY 2019 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. The ESnet-6 upgrade will also continue to identify all of the critical research necessary to deploy these technologies with continued 99.999% reliability and enhanced cyber security protections.

<u>Exascale Computing</u>
Exascale computing is a central component of a 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 ASCR FY 2019 Request includes $472,706,000 for SC's contribution to DOE's Exascale Computing Initiative to support the development of an exascale computing software ecosystem, prepare mission critical applications to address the challenges of exascale, and deploy at least one Exascale system in 2021 to meet national needs.

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, the investments in ECI will impact computing at all scales, ranging from the largest scientific computers and data centers to Department-scale computing to home computers and laptops.

Partnerships with U. S. vendors will continue to build on previous DOE investments in the Design Forward and Fast Forward pre-exascale activities to provide opportunities to accelerate our exascale efforts to keep pace with foreign competition. However, significant risk remains in delivering on this potential. To reduce this risk, the Department will increase the site preparations and NRE investments, initiated in FY 2018, between LCFs and U.S. vendors. As shown in the following table, the $472,706,000 funds continue application, software, and hardware development in SC-ECP and the site preparations and NRE activities at the LCFs to support the deployment of an exascale computing system in 2021 at ANL, followed by a second exascale system with a different advanced architecture at ORNL:

- $232,706,000 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.
- $240,000,000 in LCF activity to support operations of Theta and NRE, and site preparations for deployment of an exascale system to be delivered to the ALCF in 2021 and an additional exascale system with a different architecture at Oak Ridge in the 2021-2022 timeframe. 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 2019 Crosscuts ($K)

	ECI
Advanced Scientific Computing Research	472,706

Advanced Scientific Computing Research
Funding ($K)

	FY 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
Mathematical, Computational, and Computer Sciences Research				
Applied Mathematics	29,229	–	40,316	+11,087
Computer Science	29,296	–	38,296	+9,000
Computational Partnerships	32,596	–	62,667	+30,071
Next Generation Networking for Science	16,000	–	–	-16,000
SBIR/STTR	6,369	–	5,352	-1,017
Total, Mathematical, Computational, and Computer Sciences Research	**113,490**	–	**146,631**	**+33,141**
High Performance Computing and Network Facilities				
High Performance Production Computing	92,145	–	80,000	-12,145
Leadership Computing Facilities	190,000	–	340,000	+150,000
Exascale (non-add)	–	–	*(240,000)*	*(+240,000)*
Research and Evaluation Prototypes	26,293	–	24,452	-1,841
High Performance Network Facilities and Testbeds	45,000	–	56,435	+11,435
SBIR/STTR	16,072	–	18,786	+2,714
Total, High Performance Computing and Network Facilities	**369,510**	–	**519,673**	**+150,163**
Exascale Computing				
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	164,000	162,886	232,706	+68,706
Total, Advanced Scientific Computing Research	**647,000**	**642,606**	**899,010**	**+252,010**

SBIR/STTR funding:

- FY 2017 Enacted: SBIR $19,675,000; and STTR $2,766,000
- FY 2019 Request: SBIR $21,162,000; and STTR $2,976,000

[a] A full-year 2018 appropriation for this account was not enacted at the time the budget was prepared; therefore, the budget assumes this account is operating under the Continuing Appropriations Act, 2018 (Division D of P.L. 115-56, as amended). The amounts included for 2018 reflect the annualized level provided by the continuing resolution. (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 2019 Request vs FY 2017 Enacted
Mathematical, Computational, and Computer Sciences Research: The Computer Science and Applied Mathematics activities will increase their emphasis on data-intensive science and 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. There is a particular emphasis within the partnerships portfolio focus on future computing technologies such as quantum information science in partnership with Basic Energy Sciences (BES), Biological and Environmental Research (BER), High Energy Physics (HEP), and Nuclear Physics (NP) and on data intensive science. The Next Generation Networking for Science activity was combined with other activities in FY 2018. Collaboratory efforts previously supported by this activity are supported by computational partnerships to strengthen the interconnectivity of these efforts. Other networking-related R&D is now supported within the Computer Science activity. This funding will include new efforts in quantum networking.	**+33,141**
High Performance Computing and Network Facilities: Increased facilities funding initiates activities 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 be deployed in the 2021-2022 timeframe. Funding also supports operations, including increased power costs, upgrading, planning, and long lead site preparations at ASCR's facilities. High Performance Network Facilities and Testbeds will support the ESnet-6 upgrade to significantly increase capacity and security to all sites.	**+150,163**
Exascale Computing: The FY 2019 Request will support efforts in SC-ECP for the continued acceleration of 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.	**+68,706**
Total, Advanced Scientific Computing Research	**+252,010**

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. New partnerships with the other SC Programs have been established in quantum information 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 2019, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Leadership Computing Adds Predictive Power to New Nuclear Reactor. The largest time-dependent simulation of a nuclear power plant was run on the Oak Ridge Leadership Computing Facility to obtain accurate design predictions for the startup of Watts Bar Unit 2 nuclear power plant, the United States' first new nuclear reactor in 20 years. The simulation used data supplied by two Consortium for Advanced Simulation of Light Water Reactors (CASL) members—the Tennessee Valley Authority and the Westinghouse Electric Company. The simulations provided a detailed picture of the reactor's hour-by-hour behavior during power escalation when it came fully online in October 2016. In total, the CASL team calculated 4,128 state points, a task that required more than 2 million core-hours of compute time.

Finding New Ways to Remove Hazardous Chemicals from Mixtures. Removing hydrogen sulfide from natural gas is costly and energy-intensive, limiting the economical use of raw natural gas reserves. Researchers from the University of Minnesota and the ALCF developed new models to identify methods to remove hazardous chemicals from mixtures economically and without associated health risks. Using Mira, the researchers identified siliceous zeolites as the best candidates for removing hydrogen sulfide from mixtures of methane, ethanol, carbon dioxide, and nitrogen and simulated efficient processes that use them. The findings from these studies will provide the natural-gas industry access to untapped reserves at reduced cost and environment impact, and without health risks.

Interoperable Design of Extreme-scale Application Software. The ECP Interoperable Design of Extreme-scale Application Software project released the first version of its Extreme-scale Scientific Software Development Kit (xSDK) to improve ECP developer productivity and software sustainability while ensuring continued scientific success. The xSDK toolkit provides a superior solution for application developers using libraries by enabling turnkey installation, compatible builds, and interoperability, which is especially important for multi-scale and multi-physics projects that rely upon this functionality. The current xSDK packages include four numerical libraries, two domain components, and nine others being staged as part of future releases. The explicit ECP investment in developing, adapting, and adopting new and better software practices will improve developer productivity and software sustainability at a time when such improvements are essential for transforming capabilities for new platforms, coupling multiscale and multi-physics, and improving the effectiveness of DOE's highly skilled computational scientists.

How safe is your data? ASCR-supported researchers at the University of Wisconsin-Madison analyzed how modern distributed storage systems behave in the presence of file-system faults such as data corruption and read and write errors. They characterized eight popular distributed storage systems and uncover numerous bugs related to file-system fault tolerance. According to their findings, modern distributed systems do not consistently use redundancy to recover from file-system faults thus a single file-system fault can cause catastrophic outcomes such as data loss, corruption, and unavailability. These research findings have led to numerous bug fixes by major storage vendors.

Largest ever simulation of a quantum computer. Quantum computing uses computational elements that obey quantum mechanical laws to potentially provide transformative changes in computational power for certain high-impact tasks such as the simulation of highly complex physical systems for applications of strategic importance to DOE and SC. Although a working prototype of a universal quantum computer has not been realized yet, it is generally thought that a quantum computer deploying 49 qubits—a unit of quantum information—will be able to match the computing power of today's most powerful supercomputers. Both emulation and simulation are important for calibrating, validating and benchmarking emerging quantum computing hardware and architectures. Simulation of quantum circuits is a general method that also allows the inclusion of the effects of noise but such simulations can be very challenging even on today's fastest supercomputers. In April 2017, two researchers from the Swiss Federal Institute of Technology (ETH Zurich) have nevertheless successfully run the largest ever simulation of a quantum computer by using 8,192 of 9,688 Intel Xeon Phi processors on NERSC's newest supercomputer, Cori. The 45-qubit simulation was made possible by the performance boost gained through the use of LBNL's Roofline model during the optimization process. The Roofline model was developed by the collaboration of 2 SciDAC Institutes (SUPER and FASTMath) between 2013 and 2017. It has quickly become a broadly used performance modeling methodology across the DOE community and recently Intel has embraced the approach and integrated it into its production Intel® Advisor.

X-ray science breakthroughs enabled by advanced mathematical algorithms. An international team of researchers is gaining insights into the three-dimensional structure and behavior of nanoscale viruses using newly-developed computational tools. The breakthrough capitalized on new Multi-Tiered Iterative Phasing algorithms developed by LBNL's Center for Advanced Mathematics for Energy Applications. The algorithms and mathematics provide a new set of tools and building blocks for data analyses of single-particle imaging experiments performed at X-ray free electron laser facilities such as the DOE Linac Coherent Light Source at SLAC. In overcoming long-standing technical barriers, the mathematical developments have immediate practical benefit for accelerating research discoveries using the most advanced x-ray light sources.

How to handle vast amounts of simulation data. Today's supercomputers are capable of generating more data than they can write to disk. For traditional workflows that rely on exporting the full state of the simulation every few time steps for post-hoc analysis, this is a major hurdle that will become even more entrenched as the ability to export data does not keep pace with the ability to generate data on future computing platforms. That's one reason why scientists are now looking to do more of the analysis of the simulation data inside the supercomputer—a technique referred to as *in situ* analysis. In situ analysis, however, comes with its own challenges, for example, analysis calculations perform differently from the simulation calculations. Novel computer science techniques are needed to balance the workloads of the different calculations within the machine and manage the sharing of data between the types of calculations. Two in situ methods are used, often in combination: the first does some portion of the analysis calculation within the simulation routine; the second does a portion of the analysis alongside the simulation, but shares data with the simulation as needed. Challenges of the in situ workflow include determining which analysis tasks can be done within the simulation and which have to be done alongside it, and how to allocate resources to each. In situ technologies for HPC applications have developed largely through long-term ASCR support for basic computer science research and are being deployed in current HPC codes, enabling the science to continue despite the data output limitations of current and expected HPC machines, often while gaining better time resolution of the saved data, and more efficient use of computing and storage resources. For example, researchers at LBNL have developed such in situ techniques with support from the ASCR basic research program and, in collaboration with SciDAC activities, deployed these techniques for cosmology simulations. By saving only the analysis results rather than the full simulation data, the application saw a one thousand fold decrease in data volume exported from the simulation without loss of scientifically valuable information.

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

Description
The Mathematical, Computational, and Computer Sciences Research subprogram supports research activities to effectively meet the Office of Science HPC mission needs, including both data intensive and computationally intensive science. 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;
- new mathematics and software for data management and analysis, including machine learning, required for the complex data-intensive workflows that are increasingly reliant on HPC capabilities.
- operating systems, data management, analyses, machine learning, representation model development, 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 usability 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 supports long-term, basic research to develop methods, software, and tools to use HPC systems for data-intensive and computationally intensive science at the exascale and beyond. This requires a focus on increased parallelism, data movement, usability, and machine learning as well as exploratory research to prepare for computing paradigms in post exascale era including extremely heterogeneous architectures and beyond Moore's Law technologies such as quantum, neuromorphic, and probabilistic computing, which has the potential to revolutionize scientific computing.

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 seamless integration of HPC and data and a sophisticated tool suite for data management, analysis, and curation. In FY 2019, ASCR's research program will increase its emphasis on data-intensive science, including new efforts to exploit machine learning for discovery and decision support.

<u>Applied Mathematics</u>
The Applied Mathematics activity supports basic research leading to fundamental mathematical advances and computational breakthroughs across DOE and Office of Science missions. Basic research in scalable algorithms, multiscale modeling, and efficient data analysis underpin all of DOE's computational and data-intensive science efforts. More broadly, this activity includes support for foundational research in problem formulation, multiscale modeling and coupling, mesh discretization, time integration, advanced solvers for large-scale linear and nonlinear systems of equations, methods that use asynchrony or randomness, uncertainty quantification, and optimization. Forward-looking efforts by this activity anticipate DOE mission needs from the closer coupling and integration of scientific data with advanced computing, and for enabling greater capabilities for scientific discovery, design, and decision-support.

<u>Computer Science</u>
The Computer Science activity supports long-term, basic research on the software, tools, and techniques that allow scientists to harness the potential of advanced computing and smart networking technologies and extreme-scale data, including machine learning.

ASCR Computer Science research leads the way in the R&D of the software infrastructure that is essential for the effective use of the most powerful high performance computing systems in the country, tools to manage and analyze data at scale, and cybersecurity innovation that can enable the scientific integrity of extreme scale computation, networks, and scientific data. ASCR Computer Science plays the role of reducing risk when industry will not invest in the specialized software required for future Leadership Computers. Supercomputer vendors often take software developed with ASCR Computer Science investments and integrate it with their own software.

It is widely acknowledged that the computing industry is currently entering an era of unprecedented hardware architecture diversity — a veritable explosion of technology innovation that will greatly increase the complexity of Leadership Computing. Such diversity and complexity will be accompanied by major challenges at all levels of systems software, data management and analytics tools, the cybersecurity environment, and applications software. To have Leadership Computing systems that are useful for science applications at the time of delivery, ASCR Computer Science research focuses its investments on software needed for emerging technologies long before they are in operational use. These efforts are essential to ensure DOE mission applications including data-intensive applications, are able to use commercially available HPC through the exascale era and beyond Moore's law.

<u>Computational Partnerships</u>

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. Starting 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 will support interdisciplinary teams in partnership with BES, BER, HEP, and NP to develop algorithms and applications targeted for future computing platforms, including quantum information systems.

Activities and Explanation of Changes

FY 2017 Enacted	FY 2019 Request	Explanation of Changes FY 2019 Request vs FY 2017 Enacted
Mathematical, Computational, and Computer Sciences Research $113,490,000	**$146,631,000**	**+$33,141,000**
Applied Mathematics $29,229,000	**$40,316,000**	**+$11,087,000**
Applied Mathematics continued 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 were added to the suite of tools under development for optimizing the scientific output of data-intensive programs across SC.	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. Applied Mathematics will also continue to focus on the development of adaptive algorithms and machine learning in recognition of the increased interest in these technologies across SC application areas.	The Request increases core research efforts in applied mathematics. Priorities within the portfolio are focused on addressing opportunities for applied mathematics research to improve the rigor, robustness, and reliability of machine learning for DOE mission requirements.
Computer Science $29,296,000	**$38,296,000**	**+$9,000,000**
Computer Science continued 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 also continued, and a new effort was initiated to exploit machine learning techniques to better understand data generated both by HPC simulations and SC facilities.	Computer Science will continue efforts to develop software, new programming models, new operating systems, and efforts to promote ease of use. In addition, efforts in quantum networking, transferred from the Next Generation Networking for Science activity, will continue at FY 2018 levels. There will also be an emphasis on preparing for the “extremely heterogeneous” post-exascale era.	The Request increases core research efforts in Computer Science. Priorities within the portfolio include quantum networking and future computing technologies such as quantum, neuromorphic, or probabilistic computing.

FY 2017 Enacted	FY 2019 Request	Explanation of Changes FY 2019 Request vs FY 2017 Enacted
Computational Partnerships $32,596,000	**$62,667,000**	**+$30,071,000**
The SciDAC institutes continued 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 focused on preparing applications to harness the potential of ASCR's planned upgrades to its computing facilities. Partnerships continued to explore potential impacts of future computing technologies across SC.	In addition to continued support for the SciDAC institutes and partnerships awarded in FY 2017-18, this activity will increase efforts in quantum information science in partnership with the other SC programs, and efforts to bring the power of HPC to data intensive science.	Request supports expanded basic research efforts in quantum information science in partnership with other SC Programs such as BES, BER, HEP, and NP and data intensive science efforts, including collaboratory partnerships.
Next Generation Networking for Science $16,000,000	**$0**	**-$16,000,000**
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.	The Next Generation Networking for Science activity was combined with other activities in FY 2018. Collaboratory efforts previously supported by this activity are supported by computational partnerships to strengthen the interconnectivity of these efforts. Other networking-related R&D is now supported within the Computer Science activity.	The Next Generation Networking for Science activity was combined with other activities in FY 2018.
SBIR/STTR $6,369,000	**$5,352,000**	**-$1,017,000**
In FY 2017, SBIR/STTR funding was set at 3.65% of non-capital funding.	In FY 2019, 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 Leadership Computing Facilities (LCFs) at ORNL and ANL. These computers and the other SC research facilities generate many petabytes of data each year. Moving data to 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 and other Federal agencies, 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. In addition, the REP activity supports the Computational Sciences Graduate Fellowship to prepare the next generation of computational scientists and engineers for advanced computing systems in support of DOE workforce needs.

The facilities regularly gather requirements from the other SC research programs through a robust process 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. 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 7,000 computational scientists in about 700 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.

NERSC currently operates the 30 pf Intel/Cray system (Cori, after Nobel Laureate Gerty Cori) to support the SC research community. 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 to five years.

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) Cray system (Titan) has been one of the most powerful computers in the world for scientific research since it began operations in October of 2012 and was ranked number five on the November 2017 Top 500 list, just below the most powerful supercomputers in China, Switzerland, and Japan.[a] The FY 2018 upgrade of this facility to a 200 pf IBM/NVIDIA system (Summit) challenges the leadership of 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 system (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/Cray system (Theta) to prepare their users for the ALCF-3 upgrade in 2021.

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 with 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 2018 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 enabled 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 testbeds and prototypes.

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) is the Office of Science's high performance network user facility, delivering highly-reliable data transport capabilities optimized for the requirements of large-scale science. In essence, ESnet is the circulatory system that enables the DOE science mission. ESnet currently maintains one of the fastest and most reliable science networks in the world with a 100 gigabit per second (Gbps) "backbone" network that spans the continental United States and the Atlantic Ocean. ESnet interconnects the DOE's national laboratory system, dozens of other DOE sites, and ~200 research and commercial networks around the world—enabling tens of thousands of scientists at DOE laboratories and academic institutions across the country to transfer vast data streams and access remote research resources in real-time. ESnet also supports the data transport requirements of all SC user facilities. ESnet's traffic continues to grow exponentially—roughly 66% each year since 1990—a rate more than double the commercial internet. Costs for ESnet are dominated by operations and maintenance, including continual efforts to maintain dozens of external connections, benchmark future needs, expand capacity, and respond to new requests for site access and specialized services. As a user facility, ESnet engages directly in efforts to improve end-to-end network performance between DOE facilities and U.S. universities. ESnet is recognized as a global leader in innovative network design and operations, and is heavily engaged in planning a complete upgrade of its backbone network (the ESnet-6 upgrade).

[a] http://www.top500.org/lists/2017/11/

Activities and Explanation of Changes

FY 2017 Enacted	FY 2019 Request	Explanation of Changes FY 2019 Request vs FY 2017 Enacted
High Performance Computing and Network Facilities $369,510,000	**$519,673,000**	**+$150,163,000**
High Performance Production Computing $92,145,000	**$80,000,000**	**-$12,145,000**
The NERSC-8 system continued production operations. Demand for production computing for the SC programs continued to grow along with system capability and the rapid increase in data from experiments. In FY 2017, preparation for the 2020 delivery of NERSC-9, which will provide three to five times the capacity of NERSC-8, continued.	Support will continue operations and user support at the NERSC facility–including power, space, leases and staff. Funding will also support site preparation activities for the NERSC-9 upgrade, such as increased power and cooling capacity, and NRE efforts to ensure the new computing system meets the needs of the diverse NERSC user community.	FY 2019 funding supports site preparations and NRE activities to maintain deployment of the upgrade in the 2020 timeframe.
Leadership Computing Facilities $190,000,000	**$340,000,000**	**+$150,000,000**
Operation continued at both LCF facilities while upgrades will proceeded. The OLCF began installation and testing, the new IBM hybrid supercomputer, called Summit. This upgrade will provide 200 pf of computing capability, or approximately five times the capability of the previous system, Titan. OLCF also began site preparations to enable deployment of an exascale system in the 2021-2022 timeframe. The ALCF continued to provide access to the 8.5-pf Intel Xeon interim system, called Theta, to prepare ALCF users for the proposed 2021 exascale system. The ALCF began site preparations and significant NRE efforts to deploy a novel architecture capable of delivering more than an exaflop of computing capability in 2021.	Support will continue operations and user support at the LCF facilities –including power, space, leases, and staff. Long-lead site preparations for planned upgrades, such as increased power and cooling capacity and significant NRE efforts, will also be supported. The OLCF will continue operation and allocation of Summit. In support of ECP, the OLCF will provide access to Summit for the application and software projects to scale and test their codes. The OLCF will also continue activities to enable deployment of an exascale system in the 2021-2022 timeframe under the CORAL II. The ALCF will continue operation of Theta. The ALCF will continue site preparations and significant NRE efforts to deploy a novel architecture capable of delivering more than an exaflop of computing capability in the 2021 timeframe as part of ECI.	Increase supports significant NRE efforts and site preparation activities at both LCFs to lower the risk and support the accelerated delivery of exascale capable computing systems beginning in 2021.

FY 2017 Enacted	FY 2019 Request	Explanation of Changes FY 2019 Request vs FY 2017 Enacted
	In addition, the ALCF will procure a large developmental testbed to test activities from NRE investments and to provide ECP applications and software technology projects to test their codes.	
Leadership Computing Facility at ANL: $80,000,000	$140,000,000	+$60,000,000
Leadership Computing Facility at ORNL: $110,000,000	$200,000,000	+$90,000,000
Research and Evaluation Prototypes $26,293,000	**$24,452,000**	**-$1,841,000**
Availability of experienced and knowledgeable workforce issues continued to be of vital importance to ASCR's current and planned facilities. The CSGF program plays an increasingly important role as the Exascale Initiative progresses and future computing technologies mature. Therefore, support for the CSGF within REP continued at $10,000,000 in FY 2017. Funding also established support for the future computing technology testbed focused on quantum computing.	The Request will provide continued support for the CSGF fellowship at $10,000,000 in partnership with the NNSA to increase availability of a trained workforce for exascale and beyond Moore's Law capabilities. In addition, funding will provide continued support for quantum testbed efforts to provide resources for the researchers supported through the quantum information science partnerships with the other SC programs.	Continues support for the CSGF fellowship and quantum testbed activities. The decrease reflects the completion of a data storage and memory project.
High Performance Network Facilities and Testbeds $45,000,000	**$56,435,000**	**+$11,435,000**
Funding supported operations and maintenance of the network and continued development of tools now widely deployed through the DOE and university systems in the U.S.: Science DMZ, perfSONAR, Data Transfer Nodes, and OSCARS. ESnet was last upgraded in 2010 and some technology is no longer supported by the vendor. Additional funds were used to support planning for an upgrade to ESnet-6, which will provide a network for scientific 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.	The Request will support operations of the ESnet at 99.999% reliability. In addition, funding will support the ESnet-6 upgrade to increase network capacity and modernize the network architecture.	The Request supports the continued operations of ESnet and the ESnet-6 upgrade.
SBIR/STTR $16,072,000	**$18,786,000**	**+$2,714,000**
In FY 2017, SBIR/STTR funding was set at 3.65% of non-capital funding.	In FY 2019, SBIR/STTR funding is set at 3.65% of non-capital funding.	

Description

SC and NNSA will continue to execute the Exascale Computing Initiative (ECI), which is an effort 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 Exascale Computing Project (SC-ECP) in the Exascale Computing subprogram captures the research aspects of ASCR's participation in the U. S. DOE's ECI, to ensure the hardware and software R&D, including applications software, for a exascale system is completed in time to meet the scientific and national security mission needs of the DOE in 2021. The deployment of these systems, funded under ECI, includes necessary site preparations and NRE, 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 is 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 2019 Request includes $232,706,000 for the SC-ECP. These funds will support the preparation of mission critical applications and the development of a software stack for exascale platforms. Funding will also support additional co-design centers, vendor partnerships in preparation for an exascale system deployment in 2021, and integration between SC-ECP and the LCFs. Deployment of exascale systems will be through the LCFs as part of their usual upgrade processes. $240,000,000 of ECI funding is provided in the LCF activity to continue planning, NRE activities, and site preparations for at least two exascale systems, with the first to be delivered as early as 2021.

Activities and Explanation of Changes

FY 2017 Enacted	FY 2019 Request	Explanation of Changes FY 2019 Request vs FY 2017 Enacted
17-SC-20 Office of Science Exascale Computing Project (SC-ECP) $164,000,000[a]	**$232,706,000**	**+$68,706,000**
FY 2017 funding accelerated application and software stack development in preparation for delivery of an exascale system in 2021.	Funding will continue the acceleration of application and software stack development in preparation for delivery of an exascale system in 2021.	Request supports expanded application and software efforts and accelerations of the Pathforward vendor partnerships.

[a] In addition, $150,000,000 of ECI funding is requested within the Leadership Computing Facilities activity in FY 2018 and $240,000,000 in FY 2019 to begin planning, non-recurring engineering, and site preparations for at least one exascale system to be delivered in 2021.

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 2017	FY 2018	FY 2019
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 SC'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.		

	FY 2017	FY 2018	FY 2019
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	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 discovery, multiscale modeling, uncertainty quantification, and adaptive algorithms.	Support at least two machines learning efforts in both applied mathematics and computer science.	Support at least two partnerships in quantum information science.
Result	Met	TBD	TBD
Endpoint Target	Develop and deploy high-performance computing hardware and software systems through exascale platforms		

Capital Summary ($K)

	Total	Prior Years	FY 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
Capital operating expenses						
Capital equipment	n/a	n/a	5,000	–	5,000	–

Funding Summary ($K)

	FY 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
Research	319,855	–	422,575	+102,720
Scientific user facility operations	327,145	–	476,435	+149,290
Total, Advanced Scientific Computing Research	**647,000**	**642,606**	**899,010**	**+252,010**

[a] A full-year 2018 appropriation for this account was not enacted at the time the budget was prepared; therefore, the budget assumes this account is operating under the Continuing Appropriations Act, 2018 (Division D of P.L. 115-56, as amended). The amounts included for 2018 reflect the annualized level provided by the continuing resolution. (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 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
TYPE A FACILITIES				
NERSC	**$92,145**	**-**	**$80,000**	**-12,145**
Number of Users	6,000		6,000	–
Achieved operating hours	N/A	-	N/A	N/A
Planned operating hours	8,585		8,585	–
Optimal hours	8,585	-	8,585	–
Percent optimal hours	N/A		N/A	N/A
Unscheduled downtime hours	N/A	-	N/A	N/A

[a] A full-year 2018 appropriation for this account was not enacted at the time the budget was prepared; therefore, the budget assumes this account is operating under the Continuing Appropriations Act, 2018 (Division D of P.L. 115-56, as amended). The amounts included for 2018 reflect the annualized level provided by the continuing resolution. (These amounts are shown only at the Congressional control level and above; below that level, a dash (—) is shown).

	FY 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
OLCF	**$110,000**	**-**	**$200,000**	**+90,000**
Number of Users	1,064		1,064	–
Achieved operating hours	N/A	-	N/A	N/A
Planned operating hours	7,008		7,008	–
Optimal hours	7,008	-	7,008	–
Percent optimal hours	N/A		N/A	N/A
Unscheduled downtime hours	N/A	-	N/A	N/A
ALCF	**$80,000**	**-**	**$140,000**	**+60,000**
Number of Users	1,434		1,434	–
Achieved operating hours	N/A	-	N/A	N/A
Planned operating hours	7,008		7,008	–
Optimal hours	7,008	-	7,008	–
Percent optimal hours	N/A		N/A	N/A
Unscheduled downtime hours	N/A	-	N/A	N/A
ESnet	**$45,000**	**-**	**$56,435**	**+11,435**
Number of users[b]	N/A		N/A	N/A
Achieved operating hours	N/A	-	N/A	N/A
Planned operating hours	8,760		8,760	–
Optimal hours	8,760	-	8,760	–
Percent optimal hours	N/A		N/A	N/A
Unscheduled downtime hours	N/A	-	N/A	N/A

[a] A full-year 2018 appropriation for this account was not enacted at the time the budget was prepared; therefore, the budget assumes this account is operating under the Continuing Appropriations Act, 2018 (Division D of P.L. 115-56, as amended). The amounts included for 2018 reflect the annualized level provided by the continuing resolution. (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 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
Total Facilities	**$327,145**	**-**	**$476,435**	**+149,290**
Number of Users[b]	8,498		8,498	–
Achieved operating hours	N/A	-	N/A	N/A
Planned operating hours	31,361		31,361	–
Optimal hours	31,361	-	31,361	–
Percent of optimal hours[c]	N/A		N/A	N/A
Unscheduled downtime hours	N/A	-	N/A	N/A

Scientific Employment

	FY 2017 Enacted	FY 2018 Annualized CR[a]	FY 2019 Request	FY 2019 Request vs FY 2017 Enacted
Number of permanent Ph.D.'s (FTEs)	470	–	630	+160
Number of postdoctoral associates (FTEs)	160	–	223	+63
Number of graduate students (FTEs)	360	–	545	+185
Other scientific employment (FTEs)[d]	225	–	274	+49

[a] A full-year 2018 appropriation for this account was not enacted at the time the budget was prepared; therefore, the budget assumes this account is operating under the Continuing Appropriations Act, 2018 (Division D of P.L. 115-56, as amended). The amounts included for 2018 reflect the annualized level provided by the continuing resolution. (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 2018 PDS and does not include a new start for FY 2019.

The FY 2019 Request for SC-ECP is $232,706,000 and is an increase of $36,126,000 over the FY 2018 Request. The FY 2019 Request supports investments in application development, software technology and hardware and integration focus areas to create an exascale eco-system that supports the delivery of the first exascale-capable system in 2021 timeframe. In addition, the preliminary estimate for the total project cost was revised from $1,153,524 to $1,233,965, an increase of $80,441, which is based on updated cost information from selected application, software and hardware activities selected to participate in the project. The most recent DOE Order 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 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, SC initiated the Office of Science Exascale Computing Project (SC-ECP) within Advanced Scientific Computing Research (ASCR) to support a large research and development (R&D) co-design project between domain scientists, application and system software developers and hardware vendors to develop a exascale ecosystem as part of the ECI. Other activities included in the ECI but not the SC-ECP include $240,000,000 to support the initiation of planning, site preparations, and non-recurring engineering at both the Argonne and Oak Ridge Leadership Computing Facilities (LCFs) where the exascale machines will be housed and operated. Moreover, the LCF ECI funding will accelerate delivery of at least one exascale-capable system in 2021 timeframe. Supporting parallel development at both LCFs will reduce the overall risk of ECI and broaden the range of applications able to utilize this new capability. Procurement of exascale systems, which is not included in the ECP, will be funded within the ASCR facility budgets in the outyears. This PDS is for the SC-ECP only; prior-year activities related to the SC-ECP are also included.

In FY 2019, SC-ECP funding will support project management; continue the co-design activities between application, software, and hardware technologies; R&D of exascale systems software and tools needed for exascale programming; increased engagement and integration between SC-ECP and the LCF's upgrade projects to provide continuous testing of the ECP funded applications and software; 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	2Q FY 2019	01/03/2017	4Q FY 2019	3Q FY 2019	4Q FY 2019	N/A	4Q FY 2023
FY 2019	07/28/2016	2Q FY 2019	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:

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 and Integration:* The Hardware and Integration focus area supports vendor-based research and the integrated deployment of specific ECP application milestones and software products on targeted systems at computing facilities, including the completion of PathForward projects transitioning to facility non-recurring engineering (where appropriate), and the integration of software and applications on pre-exascale and exascale system resources at facilities.
- *Software Technology:* The 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 will deliver a high quality, sustainable product suite.
- *Application Development:* The Application Development focus area supports co-design activities between DOE mission critical applications and the software and hardware technology focus areas to address the exascale challenges: extreme parallelism, reliability and resiliency, deep hierarchies of hardware processors and memory, scaling to larger systems, and data-intensive science. As a result of these efforts, a wide range of applications will be ready to effectively use the exascale systems deployed in the 2021 timeframe under ECI.

The SC-ECP will be managed in accordance with the principles of DOE Order 413.3B, *Program and Project Management for the Acquisition of Capital Assets*, which SC uses 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

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

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 on additional hardware technologies investments, and these costs will be included in the Total Estimated Costs for the project.

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 2020– FY 2023	426,735	426,735	426,735
Total, TEC	426,735	426,735	426,735
Other project costs (OPC)			
(Research for Application Development, System Software Technology, and Hardware Technology)			
FY 2016[a]	157,944	157,944	15,615
FY 2017	164,000	164,000	100,000
FY 2018	196,580	196,580	402,909
FY 2019	232,706	232,706	232,706
FY 2020 – FY 2023	56,000	56,000	56,000
Total, OPC	807,230	807,230	807,230
Total Project Costs (TPC)			
FY 2016[a]	157,944	157,944	15,615
FY 2017	164,000	164,000	100,000
FY 2018	196,580	196,580	402,909
FY 2019	232,706	232,706	232,706
FY 2020 – FY 2023	482,735	482,735	482,735
Total, TPC	1,233,965	1,233,965	1,233,965

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	255,125	225,000	N/A
Production Ready Software	157,870	86,000	N/A
Hardware Partnerships	13,740	79,000	N/A
Total, TEC	426,735	390,000	N/A

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

	(dollars in thousands)		
	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Other Project Costs (OPC) (Research)			
Planning/Project Mgmt	109,715	118,000	N/A
Application Development	295,062	269,630	N/A
Software Research	179,303	121,423	N/A
Hardware Research	223,150	254,471	N/A
Total OPC	807,230	763,524	N/A
Total, TPC	1,233,965	1,153,524	N/A

7. Schedule of Appropriation Requests

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
FY 2019	TEC	0	0	0	0	174,735	146,000	106,000	426,735
	OPC	157,944	164,000	196,580	232,706	14,000	14,000	28,000	807,230
	TPC	157,944	164,000	196,580	232,706	188,735	160,000	134,000	1,233,965

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 2019 Budget Request, $240,000,000 is included in the LCF’s at Argonne and Oak Ridge National Laboratories 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 in 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.

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