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 sophisticated 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. ASCR supports state-of-the-art capabilities that enable scientific discovery through computation. The Computer Science and Applied Mathematics activities in ASCR provide the foundation for increasing the capability of the national high performance computing (HPC) ecosystem by focusing on long-term research to develop software, algorithms, and methods that anticipate future hardware challenges and opportunities as well as science application needs. ASCR's partnerships and coordination with industry are essential to these efforts. At the same time, ASCR partners with disciplinary sciences to deliver some of the most advanced scientific computing applications in areas of strategic importance to the Office of Science (SC) and the Department of Energy (DOE). ASCR also supports world-class, open access high performance computing facilities and high performance networks for scientific research.

For over half a century, the U.S. has 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. Leadership in HPC has also played a crucial role in sustaining America's competitiveness internationally. 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] While this continues to be true, there is also a growing recognition that the nation that leads in machine learning (ML) and artificial intelligence (AI) will lead the world in developing new technologies, medicines, industries, and military capabilities. Most of the modeling and prediction necessary to produce the next generation of breakthroughs in science, energy, medicine, and national security will come not from applying traditional theory, but from employing data-driven methods at extreme scale. Today, significant investments in Asia and Europe are challenging U.S. dominance in computing and nations around the globe are enthusiastically investing in AI. The U.S. must invest in these fields that are critical to American prosperity. Public-private partnerships remain vital as we push our state-of-the-art fabrication techniques to their limit to develop an exascale-capable (one billion billion operations per second) system while simultaneously preparing for the artificial intelligence-big data surge and what follows at the end of the current technology roadmap. 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. DOE has a long history of making fundamental contributions to applied mathematics and computer science associated with strategic computing and a similar set of contributions is foreseen for ML and AI in the science domain and related investments in advanced architectures and hardware. ASCR's proposed activities are in line with the Nation's Research and Development (R&D) priority for American Leadership in AI, Quantum Information Science (QIS), and Strategic Computing.

ASCR-supported activities are entering a new paradigm driven by sharp increases in the heterogeneity and complexity of computing systems and the need to seamlessly and intelligently integrate simulation, AI, data analysis, and other tasks into coherent and usable workflows. HPC has become an essential tool for understanding complex systems in unprecedented detail; exploring systems of systems through ensembles of simulations; learning from extreme scale, complex data; and carrying out data analyses, especially when time is of the essence. These changes are being driven by enormous increases in the volume and complexity of data generated by SC user facilities–from simulations, experiments, and observations–and these new opportunities are propelled by advances already achieved through the DOE Exascale Computing Initiative (ECI). The convergence of AI technologies with these existing investments creates a powerful accelerator for progress and gives the U.S. a distinct advantage over nations with less integrated investments.

---

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

AI and ML are critical technologies in this new paradigm that are expected to be deployed at multiple stages of the scientific process using a variety of techniques. Many popular machine learning methods lack mathematical approaches to provide robustness, reliability, and transparency and so require significant domain knowledge to be effectively applied. In addition, ML/AI applications and tools are needed to extract knowledge and discovery of patterns and classification in data from large scientific datasets that span SC programs, for example, automate data collection and advanced control and supervision of experiments at light sources, neutron sources, microscopes and telescopes; predict and avoid plasma disruptions in fusion reactors; control and optimize particle accelerators and improve the detection of events; and predict bio-design and the design of complex communities. Due to its tradition of partnering with other SC programs, its history of supporting world-leading mathematics and computer science for computation and data analysis, and its support of open access HPC facilities, which are now powerful tools for data analysis, ML, as well as simulation, ASCR is uniquely positioned to support long-term research for scientific AI and ML.

Moore's Law—the historical pace of microchip innovation whereby feature sizes reduce by a factor of two approximately every two years—is nearing an end due to limits imposed by fundamental physics; feature sizes cannot shrink smaller than the size of atoms. The emerging fields of QIS—the ability to exploit intricate quantum mechanical phenomena to create fundamentally new ways of obtaining and processing information—are opening new vistas of science discovery and technology innovation. QIS is currently at the threshold of a revolution, creating opportunities and challenges for the Nation, as growing international interest and investments are starting a global quantum race. DOE envisions a future in which the cross-cutting field of QIS increasingly drives scientific frontiers and innovations toward realizing the full potential of quantum-based applications, from computing, to communication, to sensing. This will require a detailed understanding of how quantum systems behave, accurate knowledge of how to integrate the components into complex systems, and precise control of the structures and functionalities. The traditional linear model of discovery science leading to design development and commercial deployment will not meet these goals alone within an acceptable time, due to the urgency and scale of our mission. Rather, there is a need for bold approaches that better couple all elements of the technology innovation chain and combine the talents of the program offices in SC, universities, national labs, and the private sector in concerted efforts to define and construct an internationally competitive U.S. economy. In support of the National Quantum Initiative, one or more SC QIS Centers,[a] coupled with a robust core research portfolio stewarded by the individual SC programs including ASCR, will create the ecosystem across universities, national labs, and industry that is needed to foster these developments with benefits in national security, economic competitiveness, and leadership in scientific discovery.

SC and the DOE National Nuclear Security Administration (NNSA) continue to partner on the Department's ECI to overcome key exascale challenges in parallelism, energy efficiency, and reliability, leading to deployment of a diverse set of exascale systems in the calendar year 2021-2022 timeframe. The ECI's goal for an exascale-capable system is a five-fold increase in sustained performance over the Summit HPC 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, emphasizing sustained performance in science and national security mission applications and increased convergence between exascale and large-data analytic computing.

Highlights of the FY 2020 Request

The FY 2020 Request of $920,888,000 for ASCR will strengthen U.S. leadership in strategic computing, the foundations of AI, and QIS. 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 ML/AI, and future computing technologies, and maintains support for ASCR's Computational Partnerships with a focus on developing strategic partnerships in quantum computing and data intensive applications. The Request also provides strong support for ASCR user facilities operations to ensure the availability of high performance computing and networking to the scientific community and upgrades to maintain U.S. leadership in these essential areas. Increased funding supports upgrades at 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). The Request provides robust support for ECI which includes the SC-Exascale Computing Project (SC-ECP) and site preparations, testbeds, and non-recurring engineering (NRE) activities at the LCFs in support of the delivery of at least one exascale computing system in calendar year 2021.

---

[a] Recently authorized by Section 402 of the National Quantum Initiative Act, PL 115-368.

The Request provides funding to meet the baseline schedules for the OLCF-5, NERSC-9 and ALCF-3 upgrades. In addition, to ensure the rapid and agile adoption of Big Data and AI solutions, ASCR will also support the seamless integration of data and computing resources through the ESnet-6 upgrade.

Mathematical, Computational, and Computer Sciences Research

When combined with the advances of exascale computing, ML/AI can significantly improve productivity by managing complex simulations and augmenting first principle simulations with data driven predictive models. The FY 2020 Request supports foundational research to improve the robustness, reliability, and transparency of Big Data and AI technologies, uncertainty quantification, and development of software tools to tightly couple simulation, data analysis, and AI for DOE mission applications. Investments focus on areas unique to science such as the transparency and interpretability of AI and ML, uncertainty quantification, and the computer science and software infrastructure for AI and ML applications, including tools for data management. The Request also supports partnerships among computer scientists, applied mathematicians, and domain scientists to develop hybrid models where current DOE applications, which are characterized by complex, multi-scale physics as well as large-scale, multi-faceted data, are merged with AI and ML techniques - providing the combined benefits of both techniques.

Recognizing the limits of Moore's Law, ASCR began activities in FY 2017 to explore future computing technologies, such as quantum information science (QIS) and neuromorphic computing, that are not based on silicon microelectronics. In the FY 2020 Request, QIS remains a principal emphasis. ASCR will partner with SC's Basic Energy Sciences (BES) and High Energy Physics (HEP) programs to establish at least one multi-disciplinary QIS Center to promote basic research and early stage development to accelerate the advancement of QIS through vertical integration between systems and theory and hardware and software. ASCR's Quantum Testbeds activities, which provide researchers with access to novel, early-stage quantum computing resources and services, will be expanded to support partnerships with the BES Nanoscale Science Research Centers. In addition, research in quantum information networks focuses on the opportunities and challenges of transporting and storing quantum information over interconnects and networks.

The Computer Science and Applied Mathematics activities in ASCR provide the foundation for increasing the capability of the national HPC ecosystem by focusing on long-term research to develop software, algorithms, and methods that anticipate future hardware challenges and opportunities as well as science application needs. In FY 2020, these activities will continue to address the combined challenges of increasingly heterogeneous computer architectures, and the changing ways in which HPC systems are used—incorporating more data-intensive applications and greater connectivity with distributed systems and resources, such as other SC user facilities. AI and ML are key technologies in this portfolio.

The Computational Partnerships activity is primarily focused on the Scientific Discovery through Advanced Computing (SciDAC) computational partnerships, which were re-competed in FY 2017, and 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 SciDAC portfolio will continue to focus on advancing the mission critical applications of the other SC programs. The research results emerging from the ECI inform SciDAC investments, which will, whenever possible, incorporate the software, methods, and tools developed by that initiative.

The current and predicted computing needs for DOE research and applications aggregate to a need for ubiquitous computing. Computational Partnerships also supports partnerships with other SC programs to ensure the seamless integration of Big Data and AI with computing resources to support the large-scale computing and data requirements from SC user facilities as well as to prepare for future technology through investments in QIS algorithms and applications.

High Performance Computing and Network Facilities

In FY 2020, ASCR's high performance computing and high performance networking user facilities will continue to advance scientific discovery through optimal operations. The Leadership Computing Facilities (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. NERSC will provide an innovative platform to advance SC mission research. ESnet will continue to expand capacity to meet the Department's exponential growth in scientific data traffic while executing a major upgrade to the core network.

In 2020, the ALCF will finalize site preparations and complete NRE investments with the vendor in preparations for the delivery of an exascale system (the ALCF-3 upgrade) in calendar year 2021. In addition, the ALCF will continue to operate the Theta system and provide additional testbeds for testing SC-ECP applications and software technologies at scale.

The OLCF Summit system became the world's fastest supercomputer in June 2018 and will be in full operation in FY 2020. In addition to scientific modeling and simulation, Summit offers unparalleled opportunities for the integration of AI and scientific discovery, enabling researchers to apply techniques like machine learning and deep learning to problems in high energy physics, materials discovery, and other areas. ORNL will continue site preparations, such as increased power and cooling capacity, testbeds, and NRE investments for an exascale upgrade (OLCF-5) in the calendar year 2021-2022 timeframe that will be architecturally diverse from the ALCF-3 system.

NERSC will continue operations of the 30 petaflop (pf) NERSC-8 supercomputer, named Cori. To address growing demand for capacity computing to meet mission needs, the FY 2020 Request supports activities for the delivery of NERSC-9, which will have approximately three times the capacity of NERSC-8, in late calendar year 2020. The Request also supports completion of site preparation activities for the NERSC-9 upgrade, such as increased power and cooling capacity, and investments to ensure that the diverse NERSC user community is prepared to fully utilize the new computing system.

In FY 2020, ESnet will continue to provide networking connectivity for large-scale scientific data flows while modernizing the network to meet the future needs of the DOE community. The last significant upgrade of the ESnet was in calendar year 2010, and the current optical and routing equipment is at or near the end of its operational effectiveness. The forthcoming delivery of exascale machines and the dramatically accelerating data rates from many SC user facilities and research projects demand not only ever-greater network capacity and security but also new flexibility to deliver on-demand data movement. The ESnet-6 upgrade is designed to achieve these capabilities and provide DOE with a fully integrated network backbone completely under DOE control with enhanced cyber resiliency. Funding for the upgrade continues in FY 2020.

The Department recognizes the significant and sustained competition among employers for trained computational data/network professionals, and the impact of workforce needs on achievement of the accelerated timeline for the delivery of an exascale system. The Research and Evaluation Prototypes (REP) activity will continue to support, in partnership with the NNSA, the Computational Sciences Graduate Fellowship at $10,000,000. Experienced computational scientists who assist a wide range of users in taking effective advantage of DOE's 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 at the ASCR user facilities for high end computational science and engineering through facilities operations funding. In addition, the three ASCR HPC user facilities will continue to prepare their users for future architectures through the deployment of experimental testbeds.

Exascale Computing

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 2020 Request includes $463,735,000 towards SC's contribution to DOE's ECI 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 calendar year 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 and help sustain U.S. leadership in information technology.

The results of Exascale's previous investments with vendors in the Hardware and Integration focus area were evident in the vendor's responses to the CORAL (Collaboration of Oak Ridge, Argonne and Livermore) II request for proposals for the second and third exascale systems to be sited at Oak Ridge and Lawrence Livermore National Laboratories respectively. Once the exascale system vendors have been selected, the LCFs will fund NRE activities to fully realize the potential of Exascale's vendor investments.

Investments in ECI follow the project funding plan and will help to maintain U.S. leadership in HPC into the next generation of exascale computing, which is of critical strategic importance to science, engineering, and national security. The ASCR FY 2020 Request funds two components of the ECI: planning, site preparations, and NRE at the Leadership Computing Facilities (LCF) to prepare for deployment of at least one exascale system in calendar year 2021, and the ASCR-supported Office of Science Exascale Computing Project (SC-ECP), first proposed in the FY 2017 Request, which includes the related R&D activities required to develop exascale-capable computers. The SC-ECP focuses on three areas aimed at increasing the convergence of big compute and big data, which then creates a holistic exascale HPC ecosystem:

- *Hardware and Integration:* The goal of the Hardware and Integration focus area is to integrate the delivery of SC-ECP products on targeted systems at leading DOE computing facilities.
- *Software Technology:* The goal of the 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 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 the mission of DOE, which 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, and scaling to larger systems.

Funding for ECI ($463,735,000) continues 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 calendar year 2021 at ANL, followed by a second exascale system with a different advanced architecture at ORNL:

- A total of $188,735,000 for the ECP project for the continued preparation of applications, to develop a software stack for both exascale platforms, and to support co-design centers in preparation for exascale system deployment in calendar year 2021. The final PathForward milestones were funded in FY 2019.
- A total of $275,000,000 in LCFs activity to support operations of the ALCF's Theta system and testbeds, NRE and site preparation investments at both LCFs to prepare for the deployment of an exascale system. The first exascale system will be delivered to the ALCF in calendar year 2021 and an additional exascale system, with a different architecture, will be delivered to the OLCF in the calendar year 2021-2022 timeframe. The deployment of exascale systems to these two LCFs will occur as part of their usual upgrade processes.

This approach will reduce the project risk.

ASCR supports the following FY 2020 Administration priorities.

FY 2020 Administration Priorities

	(dollars in thousands)		
	Exascale Computing Initiative (ECI)	**Artificial Intelligence (AI)**	**Quantum Information Science (QIS)**
Advanced Scientific Computing Research	463,735	36,000	51,161

Advanced Scientific Computing Research
Funding

(dollars in thousands)

	FY 2018 Enacted	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
Mathematical, Computational, and Computer Sciences Research				
Applied Mathematics	34,104	28,206	41,500	+13,294
Computer Science	29,508	22,000	38,700	+16,700
Computational Partnerships	49,910	75,667	60,959	-14,708
SBIR/STTR	4,301	4,768	5,347	+579
Total, Mathematical, Computational, and Computer Sciences Research	**117,823**	**130,641**	**146,506**	**+15,865**
High Performance Computing and Network Facilities				
High Performance Production Computing	94,000	104,000	85,000	-19,000
Leadership Computing Facilities	272,500	339,000	360,000	+21,000
Research and Evaluation Prototypes	24,260	24,452	39,453	+15,001
High Performance Network Facilities and Testbeds	79,000	84,000	80,000	-4,000
SBIR/STTR	17,417	20,701	21,194	+493
Total, High Performance Computing and Network Facilities	**487,177**	**572,153**	**585,647**	**+13,494**
Subtotal, Advanced Scientific Computing Research	**605,000**	**702,794**	**732,153**	**+29,359**
Exascale Computing				
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	205,000	232,706	188,735	-43,971
Total, Advanced Scientific Computing Research	**810,000**	**935,500**	**920,888**	**-14,612**

SBIR/STTR funding:

- FY 2018 Enacted: SBIR $19,040,000 and STTR $2,678,000
- FY 2019 Enacted: SBIR $22,329,000 and STTR $3,140,000
- FY 2020 Request: SBIR $23,269,000 and STTR $3,272,000

Advanced Scientific Computing Research
Explanation of Major Changes

	(dollars in thousands) FY 2020 Request vs FY 2019 Enacted
Mathematical, Computational, and Computer Sciences Research The Computer Science and Applied Mathematics activities will continue to increase their emphasis on the combined challenges of increasingly heterogeneous architectures, and the changing ways in which HPC systems are used—incorporating machine learning (ML) and artificial intelligence (AI) into simulations and data intensive applications while increasing greater connectivity with distributed systems and resources including other SC user facilities. The Computational Partnerships activity will continue to infuse the latest developments in applied math and computer science, particularly in the areas of AI and ML, into the strategic applications of the SC to get the most out of the leadership computing systems. These efforts will be forward funded for two years in FY 2019. In addition, the Computational Partnerships activity will continue investments in new algorithms and applications focused on both artificial intelligence and on future computing technologies such as QIS, in partnership with BES, Biological and Environmental Research (BER), High Energy Physics (HEP), and Nuclear Physics (NP). Increases in Computer Science for quantum information networks will focus on addressing new opportunities and challenges of transporting and storing quantum information.	**+15,865**
High Performance Computing and Network Facilities Increased facilities funding continues site preparations and NRE activities to deploy an exascale system at the ALCF in calendar year 2021 and for an exascale system at the OLCF, that is architecturally distinct from the ALCF system, to be deployed in the calendar year 2021-2022 timeframe. Both facilities will provide testbed resources to the SC-ECP to test and scale application codes and continuously test and deploy software technologies. In addition, funding supports the final site and early application preparations for NERSC-9 and supports the ESnet-6 upgrade to significantly increase capacity and security at all DOE sites. Funding also supports operations, including increased power costs, equipment, staffing, planning, and long lead site preparations at ASCR's facilities.	**+13,494**
Exascale Computing The FY 2020 Request will support efforts in the SC-ECP for the continuation of co-design efforts in application and software development for both planned exascale architectures and partnerships with the ASCR facilities that are providing resources for continuous integration and testing of exascale-ready software. The decrease represents completion of ASCR supported vendor partnerships with the six computer vendors to develop critical technologies, such as interconnects, processors and memory, needed for the exascale system.	**-43,971**
Total, Advanced Scientific Computing Research	**-14,612**

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 other SC Programs have been established in QIS; and the DOE activities in AI and QIS are coordinated with other agencies through the National Science and Technology Council (NSTC). There are growing areas of collaboration in the area of data-intensive science, AI, 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 R&D Subcommittee of the 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 2020, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments
Re-establishing U.S. pre-eminence in High Performance Computing. In June 2018, the Oak Ridge Leadership Computing Facility's Summit system reclaimed the top spot for the U.S. in the global "Top 500" list of high performance computing systems and held the top spot in the November list, which included five DOE systems among the top ten. Two DOE teams shared the prestigious Gordon Bell prize in 2018 for outstanding achievement in high-performance computing using the Summit system—a seven-member team affiliated with ORNL was recognized for their paper "Attacking the Opioid Epidemic: Determining the Epistatic and Pleiotropic Genetic Architectures for Chronic Pain and Opioid Addiction," and a 12-member team affiliated with the LBNL was recognized for their paper "Exascale Deep Learning for Climate Analytics."

ECP PathForward Element Drives Translational U.S. HPC Vendor Research into Exascale Platform Offerings. As evidenced by CORAL II vendor platform offerings, ECP is successfully influencing leading U.S. computing companies to maintain focus on designing and building far more powerful and balanced computers for DOE simulation workloads in spite of a strong market pull away from serving simulation workloads. This was achieved through sustained investments and management by DOE through a partnership between NNSA/ASC and SC/ASCR in processor, memory system, and interconnect R&D (FastForward, DesignForward, and now ECP's PathForward). ECP provides a path for continued American leadership in simulation even in the face of substantial market and technical challenges.

Launching the Exascale area. ORNL researchers broke the Exascale barrier, achieving a peak throughput of 1.88 exaops with mixed precision—faster than any previously reported science application—while analyzing genomic data on the recently launched Summit supercomputer. The ORNL team achieved the feat, the equivalent to carrying out nearly two billion billion calculations per second, by using a mixture of numerical precisions. Traditionally, scientific computing has relied on double-precision floating point operations. However, interest in reduced numerical precision has grown in recent years due to breakthroughs in artificial intelligence and machine learning. The ORNL researchers were able to implement high-speed single- and half-precision operations with the comparative genomics application Combinatorial Metrics (CoMet) on Summit's state-of-the-art architecture. Doing so allowed the team to achieve more than a 25-fold code speedup compared to runs conducted on the OLCF's previous leadership-class supercomputer Titan. Exascale-level performance allowed the researchers to analyze datasets composed of millions of genomes—a size that was previously impossible—and study variations between all possible combinations of two or three alleles at a time. Scientists can use this information to uncover hidden networks of genes in plants and animals that contribute to observable traits, such as biomarkers for drought-resistance in plants or disease in humans. In this demonstration, the ORNL team was able to discover key regulatory genes in plant cell walls that could be manipulated to enhance biofuels and other bioproducts. This team is also contributing to a human health application, in partnership with Department of Veteran Affairs (VA) researchers, which is a finalist for the 2018 Gordon Bell prize.

Leveraging HPCs to advance Scientific Machine Learning. Scientific data often looks very different from the data used in artificial intelligence applications. Developing the right artificial neural network can take months of handcrafting for experts and can feel like an impossible guessing game for non-experts. To expand the benefits of deep learning for science, researchers need new tools to build high-performing neural networks that don't require specialized knowledge. By leveraging the GPU computing power of the OLCF, an ORNL team has developed an evolutionary algorithm capable of auto-generating networks quickly, in a matter of hours as opposed to the months needed using conventional methods. The research team's algorithm, called MENNDL (Multinode Evolutionary Neural Networks for Deep Learning), is designed to evaluate, evolve, and optimize neural networks for unique datasets. Scaled across the OLCF Titan's 18,688 GPUs, MENNDL can test and train thousands of potential networks for a science problem simultaneously, eliminating poor performers and averaging high performers until an optimal network emerges. The process eliminates much of the time-intensive, trial-and-error tuning traditionally required of machine learning experts. One science domain in which MENNDL is already proving its value is neutrino physics. Neutrinos, ghost-like particles that pass through your body at a rate of trillions per second, could play a major role in explaining the formation of the early universe and the nature of matter—if only scientists knew more about them. The MENNDL team is working with scientists from DOE's Fermi National Accelerator Laboratory (Fermilab) to integrate neural networks into the classification and analysis of detector data. The work could improve the efficiency of some measurements, help physicists understand how certain they can be about their analyses, and lead to new avenues of inquiry. In addition to improved physics measurements, the results could provide insight into how and why machines learn–accelerating the pace of progress in scientific applications of artificial intelligence.

First simulation of an atomic nucleus using a quantum computer. Quantum computing, in which computations are carried out using uniquely quantum mechanical properties of matter, has great promise for simulating physical systems that are not accessible to conventional supercomputers. A multidisciplinary team led by scientists at ORNL took a significant step towards realizing that promise by performing the first successful simulation of an atomic nucleus using a quantum computer. The team developed a new quantum algorithm to simulate the deuteron–a proton bound to a neutron –and ran the algorithm on two different quantum computers, performing over 700,000 quantum measurements in the process. The algorithm is expected to scale up to larger and more complex atomic nuclei. This exciting new result was achieved by a collaboration between the Quantum Algorithms and Quantum Testbeds teams funded by ASCR in FY 2017 as well as the NUCLEI Scientific Discovery through Advanced Computing (SciDAC)-4 team, and leverages decades of built-up expertise in nuclear physics at ORNL and the University of Washington.

Accelerating Discovery of New Materials for dye-sensitized solar cells. Buildings consume an estimated 40 percent of energy used in the United States—a burden that also represents a renewable energy opportunity. Solar-powered windows, equipped with dye-sensitized solar cells, provide an innovative technology for generating electricity in a sustainable, environmentally friendly fashion. This type of solar cell is a promising alternative to today's solar cells made with rare earth metals that are scarce, cost prohibitive, and not environmentally sustainable. A research team from the University of Cambridge and Argonne National Laboratory is using supercomputers at the Argonne Leadership Computing Facility to combine advanced data mining techniques with machine learning and computational modeling to identify new materials with optimal properties for dye-sensitized solar cells. This approach has allowed the researchers to narrow a list of 9,000 potential materials down to six promising candidates. The team has collaborated with chemists from around the world to synthesize the six dye materials. The researchers are now working to optimize experimental conditions for solar cell device fabrication and testing. Initial experiments with the synthesized materials have yielded promising photovoltaic properties.

Convergence of big data with big compute to understand how the universe operates. NOvA, the world's longest-baseline neutrino experiment, was designed to discover more about neutrinos, ghostly yet abundant particles that travel through matter mostly without leaving a trace, to answer questions about how the universe operates. NOvA, in partnership with SciDAC, HEPcloud, and the National Energy Research Scientific Computing Center (NERSC) used over 35 million computing cores, or CPUs over approximately 54 hours to conduct the largest-scale analysis ever to support the recent evidence of antineutrino oscillation, a phenomenon that may hold clues to how our universe evolved. Although NERSC is located over 2,000 miles away from the NOvA experimental facility, researchers were able to enable near-real time analysis of time-sensitive science at rates 50 times faster than what was previously possible. Without the NERSC resources and SciDAC partnership, the NOvA collaboration could not have turned around results as quickly.

Data-driven visualization of large power grids. Driven by the emerging industry needs, electric utilities and grid coordination organizations are eager to seek advanced tools to assist grid operators and analysts to perform mission-critical tasks and enable them to make quick and accurate decisions. Traditionally, visualization of power grids heavily relies on human designers. Building and maintaining the visualization displays, however, is a very labor-intensive and error-prone process. Furthermore, the legacy approach restricts the visualization process to follow a limited number of pre-defined patterns created by human designers, thus hindering users' ability to discover. To overcome these shortcomings, researchers at Power Info LLC have developed a data-driven approach for visualization of large power grids with funding provided through the Small Business Innovation Research (SBIR) program. The developed data-driven visualization algorithm uses empirically or mathematically derived data to formulate visualizations on-the-fly. The resulting visual presentations emphasize what the data is rather than how the data should be presented, thus fostering comprehension and discovery. The software tool resulting from this research is now being leveraged by more than 70 utility organizations in North America and Europe. For example, Dominion Virginia Power used the tool, combining artificial intelligence and human's natural intelligence, to auto-generate a large number of high-quality visualization displays at a fraction of the traditional cost. Work that used to consume more than six months for two full-time employees was completed by a college intern within a month with no errors. The end-users concluded that the delivered solution saved operational costs and reduced human errors.

Most Detailed 3-D Map of Earth's Interior. Using advanced modeling and simulation, seismic data generated by earthquakes, and OLCF's Titan, a team led by Princeton University has created the most detailed 3-D picture of Earth's interior showing the entire globe from the surface to the core–mantle boundary, a depth of 1,800 miles. This first global seismic model where no approximations were used to simulate how seismic waves travel through the Earth marked a milestone for the seismology community. The model was created using seismic tomography, which is based on combining many seismograms. In the past, seismic tomography techniques have been limited in the amount of seismic data they could use. Traditional methods forced researchers to make approximations in their wave simulations and restrict observational data. The novel approach used in this study allowed researchers to use the entire dataset. Getting the most out of this data required a robust automated workflow. To improve data movement and flexibility on large-scale parallel computing resources, in collaboration with the OLCF staff, the team developed a superior file format called the Adaptable Seismic Data Format (ASDF) that leverages the Adaptable I/O System (ADIOS) parallel library, long supported by ASCR's computer science and SciDAC programs. As part of the OLCF's Center for Accelerated Application Readiness, the team is currently preparing to run further simulations on Summit to be able to image the entire globe from crust all the way down to Earth's center, including the core.

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 high performance computing (HPC) mission needs, including both data intensive and computationally intensive science. Computational science is central to progress at the frontiers of science and to our most challenging engineering problems. The Computer Science and Applied Mathematics activities in ASCR provide the foundation for increasing the capability of the national HPC ecosystem by focusing on long-term research to develop software, algorithms, and methods that anticipate future hardware challenges and opportunities as well as science application needs. ASCR partnerships and coordination with industry are essential to these efforts. ASCR partnerships with disciplinary science deliver some of the most advanced scientific computing applications in areas of strategic importance to SC. Scientific software often has a lifecycle that spans decades—much longer than the average HPC system. Research efforts must therefore anticipate changes in hardware as well as application needs over the long term. ASCR's partnerships with vendors and discipline sciences are critical to these efforts. Accordingly, the subprogram delivers:

- new mathematics and algorithms required to more accurately model systems involving processes taking place across a wide range of time and length scales and incorporate machine learning techniques into computational simulations;
- the software needed to support DOE mission applications, including new paradigms of data-intensive applications and machine learning, on current and increasingly more heterogeneous future systems;
- insights about computing systems and workflow performance and usability leading to more efficient and productive use of computing, storage and networking resources;
- collaboration tools and partnerships to make scientific resources readily available to scientists in university, national laboratory, and industrial settings; and
- long-term, basic research on future computing technologies with relevance to the DOE mission.

<u>Applied Mathematics</u>

The Applied Mathematics activity supports basic research leading to fundamental mathematical advances and computational breakthroughs across DOE and SC missions. Basic research in scalable algorithms, multiscale modeling, artificial intelligence, 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 machine learning, and for enabling greater capabilities for scientific discovery, design, and decision-support.

<u>Computer Science</u>

The Computer Science research program supports basic research that enables computing and networking at extreme scales and the understanding of extreme scale, or complex data from both simulations and experiments. Through the development of adaptive software tools, it aims to make high performance scientific computers and networks highly productive and efficient to solve scientific challenges while attempting to reduce domain science application complexity as much as possible. ASCR-supported activities are entering a new paradigm driven by sharp increases in the heterogeneity and complexity of computing systems and the need to seamlessly and intelligently integrate simulation, data analysis, and other tasks into coherent and usable workflows.

The Computer Science activity supports long-term, basic research on 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 does 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.

Computational Partnerships
The Computational Partnerships activity primarily 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.

The Computational Partnerships activity also supports critical partnerships in the areas of data analysis and future computing. Collaboratory and data analysis partnerships enable large distributed research teams to share data and develop tools for real-time analysis of the massive data flows from SC scientific user facilities, as well as the R&D of software to support a distributed data and computing environment. Interdisciplinary teams in partnership with BES, BER, HEP, and NP enable development of new algorithms and applications targeted for future computing platforms, including quantum information systems.

Activities and Explanation of Changes

FY 2019 Enacted	FY 2020 Request	Explanation of Changes FY 2020 Request vs FY 2019 Enacted
Mathematical, Computational, and Computer Sciences Research $130,641,000	**$146,506,000**	**+$15,865,000**
Applied Mathematics $28,206,000	$41,500,000	+$13,294,000
Applied Mathematics continues its core programs in new algorithmic techniques and strategies that extract scientific advances and engineering insights from massive data for DOE missions. Applied Mathematics also continues to focus on the development of adaptive algorithms and machine learning in recognition of the increased interest in these technologies across SC application areas.	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 increase investments in research to develop foundational capabilities in scientific AI and ML.	Support for the core programs continues with an increased focus on investments in the mathematical foundational of AI, such as uncertainty quantification and optimization needed to develop reliable predictive models.
Computer Science $22,000,000	$38,700,000	+$16,700,000
Computer Science continues efforts to develop software, new programming models, new operating systems, and continued efforts to promote ease of use. In addition, efforts in quantum networking, transferred from the Next Generation Networking for Science activity, continue, sustained at FY 2018 levels. In addition, there is also an emphasis on preparing for the "extremely heterogeneous" post-exascale era.	Computer Science will continue to address the combined challenges of increasingly heterogeneous architecture, and the changing ways in which HPC systems are used—incorporating more data intensive applications and greater connectivity with distributed systems and resources including other Office of Science user facilities. The Request expands efforts in quantum networking.	Support for the core program continues with an increased focus on incorporating AI and ML learning into data analytics software from networking to HPCs. Increases also support expansion of the quantum networking activity.
Computational Partnerships $75,667,000	$60,959,000	-$14,708,000
In addition to continued support for the SciDAC institutes and partnerships awarded in FY 2017-FY 2018, this activity increases efforts in QIS in partnership with the other SC programs, and efforts to bring the power of HPC to data intensive science.	In addition to continued support for the SciDAC institutes and partnerships awarded in FY 2017-18, this activity will maintain efforts in QIS in partnership with the other SC programs, and efforts to bring the power of HPC to data intensive science.	Two year AI partnerships with other Office of Science programs were forward funded in FY 2019; therefore, funding is not needed for this effort in FY 2020. The SciDAC Institutes will be recompeted and refocused on introducing new AI and ML algorithms and tools. Quantum algorithm partnerships will be re-competed as part of the proposed QIS Centers.

FY 2019 Enacted		FY 2020 Request		Explanation of Changes FY 2020 Request vs FY 2019 Enacted	
SBIR/STTR	$4,768,000		$5,347,000		+$579,000
In FY 2019, SBIR/STTR funding is set at 3.65% of non-capital funding.		In FY 2020, SBIR/STTR funding is set at 3.65% of non-capital funding.		Funding changes are the direct result of increases in the non-capital budget request.	

Advanced Scientific Computing Research
High Performance Computing and Network Facilities

Description

The High Performance Computing and Network Facilities subprogram supports the operations of forefront computational and networking user facilities to meet critical mission needs. ASCR operates three high performance computing (HPC) user facilities: the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) provides high performance computing resources and large-scale storage to a broad range of SC researchers; and the two Leadership Computing Facilities (LCFs) at ORNL and ANL provide leading-edge high performance computing capability to the U.S. research and industrial communities. ASCR's high performance network user facility, ESnet, delivers highly reliable data transport capabilities optimized for the requirements of large-scale science. Finally, operations of these facilities also includes investments in upgrades, including electrical and mechanical system enhancements, to ensure each remains state-of-the-art and can install future systems.

The Research and Evaluation Prototypes (REP) activity investigates next-generation computing systems. By actively partnering with the research community, including industry and other Federal agencies, to explore next-generation computing platforms, ASCR ensures they will serve the needs of the scientific community. Conversely, the REP activity prepares researchers to effectively use future computing platforms. Through these efforts the REP activity mitigates strategic risk. In addition, the REP activity supports DOE workforce needs through the Computational Sciences Graduate Fellowship, which prepares the next generation of computational scientists and engineers to work on advanced computing systems.

ASCR regularly gathers requirements from the other SC research programs through formal processes to inform upgrade plans. These requirements activities are also vital to planning for SciDAC and other ASCR research 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.

<u>High Performance Production Computing</u>

This activity supports NERSC at LBNL to deliver high-end production computing services for the SC research community. Approximately 7,000 computational scientists conducting 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 energy, 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 population ranges from experienced to neophyte. NERSC aids users entering the HPC arena for the first time, as well as those preparing leading-edge codes that harness the full potential of the machine.

NERSC currently operates the 30 pf Intel/Cray system (Cori). 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.

<u>Leadership Computing Facilities</u>

The LCFs enable open scientific applications, including industry applications, to harness the potential of leadership computing to advance science and engineering. The success of this effort is built on the gains made in REP and ASCR research efforts. Another LCF strength is the staff, who 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. LCF staff experience is critical to the success of industry partnerships to address the challenges of next-generation computing.

The Oak Ridge Leadership Computing Facility (OLCF) at ORNL currently operates testbeds in support of ECI and the 200 pf IBM/NVIDIA system (Summit), which achieved the global number one ranking as the world's fastest system in June 2018.

Early science applications at Summit include: 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. 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 U.S. companies. Also, the OLCF is preparing to deploy an exascale system in the calendar year 2021-2022 timeframe.

The Argonne Leadership Computing Facility (ALCF) at ANL operates an 8.5 pf Intel/Cray system (Theta) and testbeds to prepare their users and SC-ECP applications and software technology for the ALCF-3 upgrade in calendar year 2021. The ALCF-3 system, which will be the Department's first exascale system when deployed in calendar year 2021, is being designed to support the largest-scale computational simulations possible as well as large-scale analytics and machine learning. The ALCF and OLCF systems are architecturally distinct, consistent with DOE's strategy to foster diverse capabilities that provide 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 more than a factor of two.

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 emulators.

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 DOE's national laboratory system, dozens of other DOE sites, and approximately 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).

Activities and Explanation of Changes

FY 2019 Enacted		FY 2020 Request	Explanation of Changes FY 2020 Request vs FY 2019 Enacted
High Performance Computing and Network Facilities	**$572,153,000**	**$585,647,000**	**+$13,494,000**
High Performance Production Computing	$104,000,000	$85,000,000	-$19,000,000
Support continues for operations and user support at the NERSC facility–including power, space, leases and staff. Funding also supports 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.		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.	Decrease reflects the completion of many site preparation activities for deployment of NERSC-9 in early FY 2021.
Leadership Computing Facilities	$339,000,000	$360,000,000	+$21,000,000
ANL Leadership Computing Facility	*$140,000,000*	*$150,000,000*	*+$10,000,000*
ORNL Leadership Computing Facility	*$199,000,000*	*$210,000,000*	*+$11,000,000*
Support continues for 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, are supported.		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 and testbeds, will also be supported.	Support is increased for final site preparations inside the computer facilities for the planned deployment of exascale systems and for the final NRE efforts.
The OLCF continues the operation and allocation of Summit while decommissioning Titan. In support of ECP, the OLCF provides access to Summit for the application and software projects to scale and test their codes. The OLCF also continues activities to enable deployment of an exascale system in the calendar year 2021-2022 timeframe under the CORAL II.		The OLCF will continue operation and allocation of Summit. In support of ECP, the OLCF will provide access to Summit and other testbeds 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 calendar year 2021-2022 timeframe under the CORAL II.	

FY 2019 Enacted	FY 2020 Request	Explanation of Changes FY 2020 Request vs FY 2019 Enacted
The ALCF continues the operation of Theta. The ALCF continues 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. In addition, the ALCF is procuring a large developmental testbed to test activities from NRE investments and to provide ECP applications and software technology projects to test their codes.	The ALCF will continue operation of Theta and testbeds that will be deployed in FY 2019 to support SC-ECP. 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 calendar year 2021 timeframe as part of ECI, while decommissioning MIRA.	
Research and Evaluation Prototypes $24,452,000	$39,453,000	+$15,001,000
The Enacted budget provides 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 provides continued support for quantum testbed efforts to provide resources for the researchers supported through the quantum information science partnerships with the other SC programs.	The Request will maintain 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.	Increase supports new at least one new QIS center in partnership with BES and HEP.
High Performance Network Facilities and Testbeds (ESnet) $84,000,000	$80,000,000	-$4,000,000
The Enacted budget supports operations of the ESnet at 99.999% reliability. In addition, funding supports the ESnet-6 upgrade to increase network capacity and modernize the network architecture.	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 decrease reflects the execution of several significant long-lead procurement contracts for the ESnet-6 upgrade in FY 2019 and therefore fewer funding requirements in FY 2020.
SBIR/STTR $20,701,000	$21,194,000	+$493,000
In FY 2019, SBIR/STTR funding is set at 3.65% of non-capital funding.	In FY 2020, SBIR/STTR funding is set at 3.65% of non-capital funding.	Funding changes are the direct result of increases in the non-capital budget request.

Advanced Scientific Computing Research
Exascale Computing

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 Office of Science Exascale Computing Project (SC-ECP) captures the research aspects of ASCR's participation in the ECI, to ensure the hardware and software R&D, including applications software, for an exascale system is completed in time to meet the scientific and national security mission needs of DOE in calendar year 2021. The deployment of these systems, funded under ECI, includes necessary site preparations and NRE at the Leadership Computing Facilities that will ultimately house and operate the exascale systems. The ECI will execute a program, jointly 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, 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 ANL and ORNL and NERSC at LBNL.

Overall project management for the 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 projects, such as the Spallation Neutron Source project. A Memorandum of Agreement is in place between the six DOE national laboratories participating in the SC-ECP: LBNL, ORNL, ANL, Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL). The Project Office at ORNL is executing the project and coordinating among partners.

The FY 2020 Request includes $188,735,000 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. Funding for the final vendor PathForward milestones ended in FY 2019. The results of the PathForward investments were evident in the vendor's responses to the CORAL II request for proposals. Once the exascale system vendors have been selected, the LCFs will increase investments through NRE funding to fully realize the potential of the PathForward investments. Thus, the PathForward investments will no longer be needed. Deployment of exascale systems will be through the LCFs as part of their usual upgrade processes.

Activities and Explanation of Changes

FY 2019 Enacted		FY 2020 Request	Explanation of Changes FY 2020 Request vs FY 2019 Enacted
Construction	**$232,706,000**[a]	**$188,735,000**	**-$43,971,000**
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	$232,706,000	$188,735,000	-$43,971,000
Funding continues the acceleration of application and software stack development in preparation for delivery of an exascale system in 2021.		The Request will provide funding for the acceleration of application and software stack development in preparation for delivery of the first exascale system in calendar year 2021.	Decreases in funding represent completion of ASCR supported vendor partnerships while continuing investments in applications and partnerships with ASCR's facilities to continuously develop, integrate and test the ECI software stack in preparation for the delivery of the two exascale systems in the calendar year 2021–2022 timeframe.

[a] In addition, $240,000,000 of ECI funding is provided within the Leadership Computing Facilities activity in FY 2019 and $275,000,000 is requested in FY 2020 to begin planning, non-recurring engineering, and site preparations for at least one exascale system to be delivered in calendar year 2021.

Advanced Scientific Computing Research
Capital Summary

(dollars in thousands)

	Total	Prior Years	FY 2018 Enacted	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
Capital Operating Expenses Summary						
Capital equipment	N/A	N/A	10,000	5,000	5,000	—
Total, Capital Operating Expenses	**N/A**	**N/A**	**10,000**	**5,000**	**5,000**	**—**

Capital Equipment

(dollars in thousands)

	Total	Prior Years	FY 2018 Enacted	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
Capital Equipment						
Total, Non-MIE Capital Equipment	N/A	N/A	10,000	5,000	5,000	—
Total, Capital Equipment	**N/A**	**N/A**	**10,000**	**5,000**	**5,000**	**—**

Funding Summary

(dollars in thousands)

	FY 2018 Enacted	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
Research	364,500	408,500	395,888	-12,612
Facility operations	445,500	527,000	525,000	-2,000
Total, Advanced Scientific Computing Research	**810,000**	**935,500**	**920,888**	**-14,612**

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

Definitions for TYPE A facilities:
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.

(dollars in thousands)

	FY 2018 Enacted	FY 2018 Current	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
TYPE A FACILITIES					
NERSC	**$94,000**	**$94,000**	**$104,000**	**$85,000**	**-$19,000**
Number of Users	6,000	7,449	7,500	7,500	—
Achieved operating hours	N/A	8,485	N/A	N/A	N/A
Planned operating hours	8,585	8,585	8,585	8,585	—
Optimal hours	8,585	8,585	8,585	8,585	—
Percent optimal hours	99%	99%	N/A	N/A	N/A
Unscheduled downtime hours	1%	1%	N/A	N/A	N/A

	(dollars in thousands)				
	FY 2018 Enacted	FY 2018 Current	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
OLCF	**$162,500**	**$162,500**	**$199,000**	**$210,000**	**+$11,000**
Number of Users	1,064	1,444	1,450	1,450	—
Achieved operating hours	N/A	6,896	N/A	N/A	N/A
Planned operating hours	7,008	7,008	7,008	7,008	—
Optimal hours	7,008	7,008	7,008	7,008	—
Percent optimal hours	>99%	98%	N/A	N/A	N/A
Unscheduled downtime hours	<1%	2%	N/A	N/A	N/A
ALCF	**$110,000**	**$110,000**	**$140,000**	**$150,000**	**+$10,000**
Number of Users	1,434	954	950	950	—
Achieved operating hours	N/A	6,980	N/A	N/A	N/A
Planned operating hours	7,008	7,008	7,008	7,008	—
Optimal hours	7,008	7,008	7,008	7,008	—
Percent optimal hours	>99%	99%	N/A	N/A	N/A
Unscheduled downtime hours	<1%	1%	N/A	N/A	N/A
ESnet	**$79,000**	**$79,000**	**$84,000**	**$80,000**	**-$4,000**
Number of users[a]	N/A	N/A	N/A	N/A	N/A
Achieved operating hours	N/A	N/A	N/A	N/A	N/A
Planned operating hours	8,760	8,760	8,760	8,760	—
Optimal hours	8,760	8,760	8,760	8,760	—
Percent optimal hours	100%	100%	N/A	N/A	N/A
Unscheduled downtime hours	0%	0%	N/A	N/A	N/A
Total Facilities	**$445,500**	**$445,500**	**$527,000**	**$525,000**	**-$2,000**
Number of Users[b]	8,498	9,847	9,900	9,900	—
Achieved operating hours	N/A	31,121	N/A	N/A	N/A
Planned operating hours	31,361	31,361	31,361	31,361	—
Optimal hours	31,361	31,361	31,361	31,361	—
Percent of optimal hours[a]	99%	99%	N/A	N/A	N/A
Unscheduled downtime hours	1%	1%	N/A	N/A	N/A

[a] ESnet is a high performance scientific network connecting DOE facilities to researchers around the world; user statistics are not collected.
[b] Total users only for NERSC, OLCF, and ALCF.

Advanced Scientific Computing Research
Scientific Employment

	FY 2018 Enacted	FY 2019 Enacted	FY 2020 Request	FY 2020 Request vs FY 2019 Enacted
Number of permanent Ph.D.'s (FTEs)	620	611	615	+4
Number of postdoctoral associates (FTEs)	205	198	202	+4
Number of graduate students (FTEs)	516	487	495	+8
Other scientific employment (FTEs)[b]	268	256	263	+7

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

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

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

1. Summary, Significant Changes, and Schedule and Cost History

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 an exascale ecosystem as part of the ECI. Other activities included in the ECI but not the SC-ECP include $275,000,000 in FY 2020 to support the initiation of planning, site preparations, and non-recurring engineering (NRE) 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 the calendar year 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 SC-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 2020, SC-ECP funding will support project management; co-design activities between application, software, and hardware technologies; investments on critical hardware technologies with vendors, R&D of exascale systems, software, and tools needed for exascale programming; increased engagement and integration between SC-ECP and the LCF's upgrades to provide continuous integration and testing of the ECP funded applications and software; and completion of the milestones of the vendor partnerships which received final funding in FY 2019.

Significant Changes

This Project Data Sheet (PDS) is an update of the FY 2019 PDS and does not include a new start for FY 2020.

The FY 2020 Request for SC-ECP is $188,735,000 and is a decrease of $43,971,000 from the FY 2019 Enacted. The FY 2020 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 the calendar year 2021 timeframe. The project is expected to achieve CD-2 in early FY 2020 and the decrease is a result of ASCR funding in FY 2019 its share of the vendor partnerships to initiate their final milestones

Following the Independent Project Review in January 2018, a Baseline Change Proposal (BCP) was executed and approved in March 2018, to officially move the responsibility for the exascale systems' NRE and development testbeds from ECP scope to the Exascale Computing Initiative (ECI) through the DOE high performance computing facilities. This effectively eliminated the need for CD-3A.

The current preliminary estimate for the SC-ECP total project cost was revised from $1,233,965 to $1,256,385, an increase of $22,420, which is based on updated cost information from the remaining 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, which was approved on January 3, 2017. Even with the BCP, the estimated Total Project Cost (TPC) range of the SC-ECP is $1.0 billion to $2.7 billion.

A Federal Project Director with the appropriate certification level was assigned to this project.

Critical Milestone History

Fiscal Year	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
FY 2020	07/28/2016	2Q FY 2019	01/03/2017	1Q FY 2020	3Q FY 2019	1Q FY 2020	N/A	4Q FY 2023

CD-0 – Approve Mission Need for a construction project with a conceptual scope and cost range
Conceptual Design Complete – Actual date the conceptual design was or will be completed (if applicable)
CD-1 – Approve Alternative Selection and Cost Range
CD-2 – Approve Performance Baseline
Final Design Complete – Estimated/Actual date the project design will be/was complete(d)
CD-3A – Approve Long Lead Time Procurements
CD-3B– Approve Start of Construction
D&D Complete – Completion of D&D work
CD-4 – Approve Start of Operations or Project Closeout

Fiscal Year	Performance Baseline Validation
FY 2017	TBD
FY 2018	4Q FY 2019
FY 2019	4Q FY 2019
FY 2020	1Q FY 2020

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.

(dollars in thousands)

Fiscal Year	TEC, Design	TEC, Construction	TEC, Total	OPC, Except D&D	OPC, D&D	OPC, Total	Total
FY 2017	N/A	TBD	TBD	TBD	N/A	TBD	TBD
FY 2018	N/A	390,000	390,000	763,524	N/A	763,524	1,153,524
FY 2019	N/A	426,735	426,735	807,230	N/A	807,230	1,233,965
FY 2020	N/A	426,735	426,735	829,650	N/A	829,650	1,256,385

2. 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.

[a] http://science.energy.gov/ascr/research/scidac/exascale-challenges

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 calendar year timeframe under ECI.

Justification

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 the National Energy Research Scientific Computing Center (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. Once the project is baselined in FY 2020, 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.

Key Performance Parameters (KPPs)

The Threshold KPPs represent the minimum acceptable performance that the project must achieve. Achievement of the Threshold KPPs will be a prerequisite for approval of CD-4, Project Completion. The Objective KPPs represent the desired project performance.

Performance Measure	Threshold	Objective
Performance on scientific and national security applications relative to today's performance	Greater than 50x improvement in performance[a] by at least 50% of the subset of ECP applications selected for measurement under KPP-1.	Greater than 50x improvement in performance by 100% of subset of ECP applications selected for measurement under KPP-1.
Broaden the reach of exascale science and mission capability	50% of the subset of ECP applications selected for measurement under KPP-2 s can execute their exascale challenge problem[b]	100% of the subset of ECP applications selected for measurement under KPP-2 can execute their exascale challenge problem

[a] Performance is measured by a Figure of Metric that represents the rate of "science work" defined specific to each scientific application and takes into consideration the increased complexity and precision in addition to the speed of solution.

[b] This KPP assesses the successful creation of new exascale science and mission capability. An exascale challenge problem is defined for every scientific application in the project. The challenge problem is reviewed annually to ensure it remains both scientifically impactful to the nation and requires exascale-level resources to execute.

Performance Measure	Threshold	Objective
Productive and Sustainable High-Performance Computing (HPC) software ecosystem	Software teams meet 75% of their impact goals[a]	Software teams meet 100% of their impact goals
Enrich the HPC Hardware Ecosystem	Vendors meet 80% of all the PathForward milestones	Vendors meet 100% of all the PathForward milestones

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
(Hardening of Applications Development System Software Technology, Hardware Technology)			
FY 2020	174,735	174,735	174,735
Outyears	252,000	252,000	252,000
Total, TEC	**426,735**	**426,735**	**426,735**
Other Project Costs (OPC)			
(Research for Application Development, System Software Technology, and Hardware Technology)			
FY 2016[b]	157,944	157,944	8,338
FY 2017	164,000	164,000	89,058
FY 2018	205,000	205,000	180,453
FY 2019	232,706	232,706	230,066
FY 2020	14,000	14,000	250,000
Outyears	56,000	56,000	71,735
Total, OPC	**829,650**	**829,650**	**829,650**
Total Project Costs (TPC)			
FY 2016[c]	157,944	157,944	8,338
FY 2017	164,000	164,000	89,058
FY 2018	205,000	205,000	180,453
FY 2019	232,706	232,706	230,066
FY 2020	188,735	188,735	424,735
Outyears	308,000	308,000	323,735
Total, TPC	**1,256,385**	**1,256,385**	**1,256,385**

4. Details of 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.

[a] This KPP measures progress on the goal to develop a software ecosystem where high-performance applications can be efficiently and effectively designed, developed, tuned, and executed on exascale systems. Each software effort in the project defines 2-4 impact goals, which must be measurable and provide tangible value to the HPC ecosystem.

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

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Application Development	211,800	255,125	N/A
Production Ready Software	147,190	157,870	N/A
Hardware Partnerships	67,745	13,740	N/A
Total, TEC	**426,735**	**426,735**	**N/A**
Other Project Costs (OPC) (Research)			
Planning/Project Mgmt	123,715	109,715	N/A
Application Development	298,824	295,062	N/A
Software Research	167,970	179,303	N/A
Hardware Research	239,141	223,150	N/A
Total, OPC	**829,650**	**807,230**	**N/A**
Total Project Cost (TPC)	**1,256,385**	**1,233,965**	**N/A**

5. Schedule of Appropriation Requests

(dollars in thousands)

Request Year	Type	FY 2016[a]	FY 2017	FY 2018	FY 2019	FY 2020	Outyears	Total
FY 2017	TEC	—	—	TBD	TBD	TBD	TBD	TBD
	OPC	157,894	154,000	TBD	TBD	TBD	TBD	TBD
	TPC	157,894	154,000	TBD	TBD	TBD	TBD	TBD
FY 2018	TEC	—	—	—	—	175,000	215,000	390,000
	OPC	157,944	164,000	196,580	189,000	14,000	42,000	763,524
	TPC	157,944	164,000	196,580	189,000	189,000	257,000	1,153,524
FY 2019	TEC	—	—	—	—	174,735	252,000	426,735
	OPC	157,944	164,000	196,580	232,706	14,000	42,000	807,230
	TPC	157,944	164,000	196,580	232,706	188,735	294,000	1,233,965
FY 2020	TEC	—	—	—	—	174,735	252,000	426,735
	OPC	157,944	164,000	205,000	232,706	14,000	56,000	829,650
	TPC	157,944	164,000	205,000	232,706	188,735	308,000	1,256,385

6. 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 or research program's budget. A BCP was executed in March, 2018 to reflect this change. In the FY 2020 Budget Request, $275,000,000 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 (fiscal quarter or date)	2022
Expected Useful Life (number of years)	5
Expected Future start of D&D for new construction (fiscal quarter)	4Q 2030

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

7. D&D Information

N/A, no construction.

8. 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.