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, methods, tools and workflows that anticipate future hardware challenges and opportunities as well as science application needs. ASCR's partnerships and coordination with the other Office of Science (SC) programs and 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 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 regular deployment of new computing systems and advanced networks along with the applied mathematics and software technologies to effectively use 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." While this continues to be true, there is also a growing recognition that the nation that leads in artificial intelligence (AI) and machine learning (ML) and in the integration of the computing and data ecosystem 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 tightly coupled to experiments and scientific user facilities.

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 for science. 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. The FY 2021 Request strengthens partnerships with the other SC programs to build the computational and data infrastructure needed to support scientific computing from the edge to ASCR high performance computing facilities.

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, AI/ML 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, to automate data collection and advanced control and supervision of experiments at light sources, neutron sources, microscopes, and telescopes; to predict and avoid plasma disruptions in fusion reactors; to control and optimize particle accelerators and improve the detection of events; and to predict bio-design and the design of complex communities. There are significant infrastructure and integration challenges to be addressed to realize the potential of these scientific investments. Because of 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.

The emerging field within quantum information science (QIS) – the ability to exploit intricate quantum mechanical phenomena to create fundamentally new ways of obtaining and processing information – is opening new vistas of science discovery and technology innovation that build on decades of investment across SC. QIS is currently at the threshold of a potentially disruptive 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 sensing, connected through a quantum internet. 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 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 this 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, SC QIS Centers established in FY 2020 constitute an interdisciplinary partnership between ASCR and the other SC programs. This partnership complements a robust core research portfolio stewarded by the individual SC programs, including ASCR, to create the ecosystem across universities, national laboratories, and industry that is needed to advance developments in QIS and related technology.

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 and economics. As a result, numerous emerging technologies are competing to help sustain productivity gains, each with its own risks and opportunities. The challenge for ASCR is in understanding their implications for scientific computing and being ready for the potential disruptions from rapidly evolving technologies without stifling innovation or hampering scientific progress. ASCR's strategy is to focus on technologies that build on expertise and core investments across SC, continuing engagements with industry and the scientific community from the exascale computing project, investing in small-scale testbeds and increasing core research investments in Applied Mathematics and Computer Science.

Today, significant investments in Asia and Europe are challenging U.S. dominance in computing, and nations around the globe are enthusiastically investing in AI and QIS. The U.S. must invest in the fields and infrastructure 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, with its integration and infrastructure challenges, and the "Cambrian explosion" of specialized hardware expected as we reach 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 foresees making a similar set of contributions for AI/ML in the science domain and related investments in advanced architectures, scientific data infrastructure, and emerging technologies. ASCR's proposed activities are in line with the Administration's FY 2021 Research and Development (R&D) priority for American Leadership in AI, Quantum Information Science (QIS), and Strategic Computing.

SC and the National Nuclear Security Administration (NNSA) continue to partner in the Department's ECI to overcome key exascale challenges in parallelism, energy efficiency, and reliability, with emphasis on the implications for both simulation and data science at this scale, leading to deployment of a diverse set of two exascale systems by SC 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 2021 Request

The FY 2021 Request of $988,051,000 for ASCR will strengthen U.S. leadership in strategic computing, the foundations of AI and QIS, and the infrastructure that enables data-driven science. To ensure ASCR is meeting SC's HPC mission needs during and after the exascale project, the Request prioritizes basic research in Applied Mathematics and Computer Science with emphasis on the challenges of data intensive science, including AI/ML, and future computing technologies, and significantly increases support for ASCR's Computational Partnerships with a focus on developing strategic partnerships in quantum computing and data intensive applications and new partnerships that broaden the impact of both exascale and data infrastructure investments in areas of strategic importance to DOE and SC. The Request also provides strong support for ASCR user facilities operations to ensure the availability of high performance computing, data, and networking to the scientific community and upgrades to maintain U.S. leadership in these essential areas. 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 supports testbeds both at the facilities and through the Advanced Computing Research activity. The Request maintains support for the Computational Sciences Graduate Fellowship (CSGF) with increased emphasis on foundational research in AI and QIS. It 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.

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 integration of data and computing resources through the ESnet-6 upgrade and building on the research and partnership investments in data infrastructure, tools, and workflows.

<u>Research</u>

- Advances in exascale computing, AI/ML, and a robust data infrastructure when combined can significantly improve scientific productivity by managing complex simulations and augmenting first principle simulations with data driven predictive models. The 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/ML, uncertainty quantification, and the computer science and software infrastructure for AI/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/ML techniques—providing the combined benefits of both techniques. These efforts will be combined with new efforts to develop the computational and data infrastructure to more seamlessly integrate SC user facilities and data repositories with compute resources and other tools for extracting scientific insights.

- Recognizing the limits of Moore's Law, ASCR began activities in FY 2017 to explore future computing technologies, such as QIS and neuromorphic computing, that are not based on silicon microelectronics. QIS remains a principal emphasis due to the potential for disruption and the significant expertise and investments across SC. ASCR will continue to partner with the other SC programs to support multi-disciplinary QIS Centers. These centers 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 continue. In addition, basic research in quantum information networks will focus on the opportunities and challenges of transporting and storing quantum information over interconnects and networks toward a vision to deliver a fundamentally new capability. In FY 2021, ASCR will support early stage research associated with the first steps to establishing a dedicated Quantum Network.

- 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 smart software, algorithms, and methods that explore the wide array of emerging technologies and anticipate future hardware challenges and opportunities as well as science application needs. In FY 2021, these activities will be expanded to address the combined challenges of increasingly heterogeneous computer architectures, and the changing ways in which HPC systems are used –

incorporating and addressing critical data science research to support the incorporation of AI into simulation and experiments including greater connectivity with distributed systems and resources, such as other SC user facilities. Emerging AI and data infrastructure technologies are a significant focus of this portfolio.

- The Computational Partnerships activity is primarily focused on the Scientific Discovery through Advanced Computing (SciDAC) computational partnerships and Institutes to advance and apply the software, tools, and methods developed by ASCR core research efforts in Applied Mathematics and Computer Science. 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 support recently recompeted SciDAC Institutes, and will recompete the SciDAC applications with emphasis on ECP and emerging AI and data infrastructure technologies. The research results emerging from the ECI inform SciDAC investments, which will, whenever possible, incorporate the software, methods, and tools developed by that initiative. Building on the success of the ECI funded partnership with the National Cancer Institute, funding is requested to partner with the NIH to expand the capabilities of DOE's tools and address NIH's rapidly growing data and computational challenges.

- ASCR will partner with SC's Offices of Basic Energy Sciences (BES), High Energy Physics (HEP), and Fusion Energy Sciences (FES) to support multi-disciplinary microelectronics research that will promote basic research to accelerate the advancement of microelectronic technologies in a co-design innovation ecosystem in which materials, chemistries, devices, systems, architectures, algorithms, and software are developed in a closely integrated fashion.

- The current and predicted computing and data-intensive needs of DOE research and applications aggregate to a need for a robust, integrated, computing and data infrastructure. All research activities will also support 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 partnerships in emerging areas such as QIS.

- To address the need for strong connections between core research and prototype development, the Research and Evaluation Prototypes (REP) activity is transferred from the Facilities subprogram to the Research subprogram under the Advanced Computing Research activity beginning in FY 2021. This activity continues to support, in partnership with NNSA, the Computational Sciences Graduate Fellowship in FY 2021 at $10,000,000 with increased focus on AI and QIS. This activity also supports the Quantum computing testbeds and early stage research activities associated with a Quantum internet (Qnet) as well as small investments in cybersecurity and microelectronics research.

Facility Operations

- In FY 2021, 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 2021, the ALCF will install the first U.S. exascale computing system, named Aurora, and begin acceptance and early science testing. In addition, the ALCF will continue to operate the Theta system and provide additional testbeds for testing SC applications and software technologies, including AI, at scale.

- The OLCF Summit system became the world's fastest supercomputer in June 2018 and remains the fastest according to the November 2019 Top500 List[a]. Summit will remain in operation throughout FY 2021. In addition to scientific modeling and simulation, Summit offers unparalleled opportunities for the integration of AI and data intensive 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 make final site preparations and NRE investments

[a] https://www.top500.org/list/2019/11/

for an exascale upgrade (OLCF-5), named Frontier, in the calendar year 2021-2022 timeframe. Frontier will be architecturally diverse from the Aurora ALCF-3 system.

- NERSC will continue operations of the 30 petaflop (pf) NERSC-8 supercomputer, named Cori and begin operations of the 75 petaflop NERSC-9 system, named Perlmutter after LBNL Nobel Laureate Saul Perlmutter. To keep pace with the increasing demand from SC researchers for AI, simulation, and data-intensive applications on NERSC, the Request also supports planning for expanded capacity and investments to ensure that the diverse NERSC user community is prepared to utilize exascale and emerging computing systems.

- In FY 2021, 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 and data infrastructure integration to deliver on-demand data management. 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 2021.

- 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. 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, the Request continues support for post-doctoral programs 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.

Projects

- Exascale computing is a central component of a long-term collaboration between SC's ASCR program and NNSA's Advanced Simulation and Computing (ASC) program to maximize the benefits of DOE's investments, avoid duplication, and leverage the significant expertise across the DOE complex. The ASCR FY 2021 Request includes $438,945,000 for 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., DOD and NIH. The importance of exascale computing to the other SC 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 and selected architectures for the exascale systems.

- 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 2021 Request funds two components of the ECI: final 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 ($438,945,000) continues with a focus on final preparation of applications and software in SC-ECP and the final preparations at the LCFs to support the installation of Aurora at ANL, followed by Frontier at ORNL. This approach will reduce the project risk:

- A total of $168,945,000 for the SC-ECP project for the final preparation of applications and the software stack for testing on both exascale platforms, and continued support for co-design centers preparation for exascale system installation in calendar year 2021.
- A total of $270,000,000 in LCFs activity to support operations of the ALCF's Theta system and testbeds, final site preparation investments needed just prior to the installation of Aurora, and to support final site preparations at the OLCF needed just prior to installation of Frontier. The deployment of exascale systems to these two LCFs will occur as part of their usual upgrade processes and, once accepted, the systems' the lease payments will be funded as part of operations.

FY 2021 Research Initiatives

Advanced Scientific Computing Research supports the following FY 2021 Research Initiatives.

	(dollars in thousands)			
	FY 2019 Enacted	**FY 2020 Enacted**	**FY 2021 Request**	**FY 2021 Request vs FY 2020 Enacted**
New Research Initiatives				
Data and Computational Collaboration with NIH	—	—	1,000	+1,000
Integrated Computational and Data Infrastructure for Scientific Discovery	—	—	11,845	+11,845
Total, New Research Initiatives	**—**	**—**	**12,845**	**+12,845**
Ongoing Research Initiatives				
Artificial Intelligence and Machine Learning	15,000	36,000	56,000	+20,000
Exascale Computing Initiative	472,706	463,735	438,945	-24,790
Quantum Information Science	33,666	54,680	86,162	+31,482
Microelectronics Innovation	—	—	5,000	+5,000
Total, Ongoing Research Initiatives	**521,372**	**554,415**	**586,107**	**+31,692**

Advanced Scientific Computing Research
Funding

(dollars in thousands)

	FY 2019 Enacted	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
Mathematical, Computational, and Computer Sciences Research				
Applied Mathematics	28,206	41,500	53,728	+12,228
Computer Science	22,000	38,700	49,605	+10,905
Computational Partnerships	75,667	69,142	75,051	+5,909
Advanced Computing Research	—	—	76,007	+76,007
SBIR/STTR	4,768	5,658	9,637	+3,979
Total, Mathematical, Computational, and Computer Sciences Research	**130,641**	**155,000**	**264,028**	**+109,028**
High Performance Computing and Network Facilities				
High Performance Production Computing	104,000	110,000	85,000	-25,000
Leadership Computing Facilities	339,000	375,000	370,000	-5,000
Research and Evaluation Prototypes	24,452	39,000	—	-39,000
High Performance Network Facilities and Testbeds	84,000	90,000	80,000	-10,000
SBIR/STTR	20,701	22,265	20,078	-2,187
Total, High Performance Computing and Network Facilities	**572,153**	**636,265**	**555,078**	**-81,187**
Subtotal, Advanced Scientific Computing Research	**702,794**	**791,265**	**819,106**	**+27,841**
Construction				
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	232,706	188,735	168,945	-19,790
Total, Advanced Scientific Computing Research	**935,500**	**980,000**	**988,051**	**+8,051**

SBIR/STTR funding:
- FY 2019 Enacted: SBIR $22,329,000 and STTR $3,140,000
- FY 2020 Enacted: SBIR $25,160,000 and STTR $3,538,000
- FY 2021 Request: SBIR $26,051,000 and STTR $3,664,000

Advanced Scientific Computing Research
Explanation of Major Changes

	(dollars in thousands) FY 2021 Request vs FY 2020 Enacted
Mathematical, Computational, and Computer Sciences Research The Computer Science and Applied Mathematics activities will continue to increase their efforts on the combined challenges of increasingly heterogeneous architectures, and the changing ways in which HPC systems are used, incorporating AI and ML 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 data infrastructure tools, into strategic applications, including new areas such as revolutionizing microelectronics, to get the most out of the leadership computing systems and data infrastructure investments. In addition, the Computational Partnerships activity will increase investments in new algorithms, applications, and data infrastructure, focused on both artificial intelligence and on future computing technologies, such as QIS and bio-inspired/bio-accelerated computing in partnership with the other SC programs and other partners such as NIH. Increases in Computer Science for quantum information networks will focus on addressing new opportunities and challenges of transporting and storing quantum information. Within the Advanced Computing Research activity, Research and Evaluation Prototypes provides support for emerging technology testbeds with emphasis on Quantum, in close coordination with the other SC programs. This includes early stage research activities for a Quantum internet (Qnet). This activity also supports the Computational Sciences Graduate Fellowship (CSGF).	**+109,028**
High Performance Computing and Network Facilities Facilities funding is reduced due to the shift of Research and Evaluation Prototypes activities to the Research subprogram, and to the completion of significant site preparation activities at the LCFs and NERSC. The Request supports final site preparations to deploy an Aurora exascale system at the ALCF in calendar year 2021 and the architecturally distinct Frontier exascale system at the OLCF, to be deployed in the calendar year 2021-2022 timeframe. Both facilities will continue to provide testbed resources to the SC-ECP and other SC researchers to test and scale application codes and continuously test and deploy software technologies. In addition, funding supports operation of the NERSC-9 and Cori systems 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 lease payments at ASCR's facilities.	**-81,187**
Exascale Computing The FY 2021 Request will support efforts in the SC-ECP for the finalization 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 that developed critical technologies, such as interconnects, processors and memory, needed for the exascale system and accelerated the transition from research to testing of applications and software on the exascale testbeds.	**-19,790**
Total, Advanced Scientific Computing Research	**+8,051**

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 DOE. 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 2021, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Getting Ready for Exascale.

The efficiency of the Spectral Neighbor Analysis Potential (SNAP) machine-learning kernel on various architectures displayed a worrisome decrease in terms of performance relative to the theoretical peak of the hardware, particularly on GPU accelerated architectures. Addressing this concern has become a central effort of the Exascale Computing Project's (ECP) Exascale Atomistics for Accuracy, Length, and Time (EXAALT) project. A team comprised of NERSC staff, the ECP Co-design center for Particle Applications (CoPA), and the ECP software team has completely re-engineered the SNAP kernel from the ground up. Beginning with a stand-alone Compute Unified Device Architecture (CUDA) prototype, re-engineered SNAP kernel has been implemented in the widely used molecular dynamics software, Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) using an ECP library for performance portability. This implementation reversed the trend and led to a five times improvement in performance. Deploying the new SNAP kernel on Summit results in a simulation throughput of 175,100 atoms-steps/wall-clock second per node, exploiting all 6 General Processing Units (GPUs) and 42 Central Processing Unit (CPU) cores.

*HP-CONCORD Enables Machine Learning in HPC at Unprecedented S*cale.

Understanding the interactions between millions of variables such as how a disease may be caused by a subset of the human genes is among the most challenging problems in data-driven science. A powerful machine learning algorithm called CONCORD can identify these relationships however until recently it could only be applied to modest-sized data sets. Researchers from LBNL supported by SciDAC along with their collaborators from Google Brain, University of California Santa Barbara and Carnegie Mellon University developed a high-performance computing version of the algorithm called HP-CONCORD. Using NERSC, they then demonstrated this parallel algorithm on an enormous set of data from the Human Connectome Project, which computed estimates for about 4 billion parameters, and an even larger demonstration problem with over 800 billion parameters. The HP-CONCORD team used functional magnetic resonance imaging data to estimate the underlying conditional dependency structure of the brain and then used the resulting estimate to automatically identify functional regions of the brain. The researchers expect that many other science areas will benefit from HP-CONCORD in the future: In biology it can be used to reconstruct gene regulatory networks, or in environmental sciences it can help in estimating temperature-to-environmental-proxy relationships. HP-CONCORD may also be useful for hypothesis generation in exploratory data analysis to guide further experimental study.

Autonomous AI-driven approach maximizes the value of high-throughput experiments.

What can be done to help a scientist decide which sequence of experiments to run in order to get the most information as quickly as possible? Researchers in the Center for Advanced Mathematics for Energy Research Applications (CAMERA) have developed a way to optimize beamline experiments and automatically choose the best next experiment to perform. New computational methods were developed to quickly explore the parameter space of predicted physical effects in a way that reduces the overall beamline time used and increases the yield from the experiment. Actions are carried out in real-time in beamline experiments at NSLS-II (BNL) and the ALS (LBNL). At BNL, beamline utilization increased from 20 percent to 80 percent, with a six-fold decrease in the number of experiments required to achieve similar results.

Quantum Noise Mitigation for the Masses.
While quantum computers today are very noisy, much of the noise is systematic and can be characterized with good precision. A team at ORNL and Virginia Tech have adapted a standard technique in experimental physics — amplifying noise in order to treat it as a "signal" and then examine how that noise varies with respect to a degree of freedom, or "knob" that we have control over – to develop a post-processing technique to mitigate the effect of noise on the output of today's quantum processors. A U.S. company took note of the result and made a pull request to the ASCR-funded team, who responded by making their technique available as a new module in the company's widely used, open source software development kit for quantum computing.

Dark optical fiber proves to be a promising earthquake sensor.
In one of the first case studies to employ a large regional network as an earthquake sensor, researchers from Lawrence Berkeley National Laboratory used distributed acoustic sensing along a 20-mile segment of the 13,000-mile-long ESnet Dark Fiber Testbed to collect and analyze seven months of passive seismic data. Because the ESnet Testbed has regional coverage, they were able to monitor seismic activity and environmental noise with finer detail than previous studies. The research team used ambient noise interferometry techniques to process 300 terabytes of collected data and extract surface wave velocity information, which enabled them to map shallow structural profiles and groundwater depth and detect regional and teleseismic earthquakes. In traditional seismology, researchers studying how the earth moves in the moments before, during, and after an earthquake rely on sensors that cost tens of thousands of dollars to make and install underground. This new approach could augment the performance of earthquake early warning systems currently being developed in the Western U.S., with higher sensitivity and at lower cost.

Hackathon Leads to Game-Changing Code Speedup for General Electric (GE).
GE Power technologies generate approximately one-third of the world's electricity. In 2018, Hyperion Research published a case study on how GE engineers were able to nearly double the efficiency of their gas turbines with the help of supercomputing simulation[a]. Recently, GE has become increasingly interested in improving GPU performance to harness the power of new HPC architectures to run its codes faster to create more design iterations and better end designs. A team from GE participated in two of the OLCF's 2019 hackathons—the first at Brookhaven National Laboratory last September and the second at the Massachusetts Institute of Technology—and used the CUDA programming model on the Summit supercomputer to gain 50-fold to 300-fold speedups in portions of their computational fluid dynamics code GEneralized Nonlinear Extension of Surface Integral Schemes (GENESIS). GENESIS can help engineers predict turbulence in complex geometries, such as those in multiple rows of turbine blades. These predictions can lead to more efficient gas turbine engines, increasing their competitiveness. The GE team has taken a typical month-long simulation down to just over two hours and has gained the ability to simulate several rows of turbine blades rather than just a single blade, demonstrating a leap from component-level to system-level simulations. Additionally, the simulations have better predictive accuracy and higher resolution than any of the company's previous simulations. The value gained from running on Summit's GPUs has also enabled the team to plan for GE's next high-performance computing investment.

New biofuel process is identified through HPC screening.
Separation of ethanol from water, essential for biofuel production, is usually effected via distillation. A new, patented process—by which ethanol would be preferentially adsorbed from a water mixture through a zeolite-containing adsorbent or membrane—would permit significant energy savings over traditional distillation-based separation. The material for the membrane was identified by a team of researchers from the University of Minnesota who used ALCF resources to conduct high throughput computational screening to identify promising microporous zeolites for various biofuel production processes. The team scaled their codes to more than 100,000 cores and simulate approximately 400 zeolites over a range of solution-phase compositions (ranging from 0.1 to 40 percent ethanol by weight). Each candidate zeolite required numerous simulations so that meaningful statistics could be collected and subsequently averaged to attain a higher degree of fidelity. Without access to supercomputer and a large allocation of computer time, it would have taken much longer to examine far fewer zeolites and the best performing ones may not have been found.

---

[a] https://insidehpc.com/white-paper/understanding-behaviors-extreme-environment-natural-gas-turbine-generators/

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 SC high performance computing (HPC) mission needs, including both data intensive and computationally intensive science. Computational and data intensive sciences are 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 and data infrastructure 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's 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 and rapidly developing capabilities such as AI and QIS, 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 AI and ML techniques into HPC simulations;
- the software needed to support DOE mission applications, including new paradigms of data-intensive applications, AI and scientific 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, data infrastructure and partnerships to make scientific resources readily available to scientists in university, national laboratory, and industrial settings;
- expertise in applying new algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to SC and DOE; and
- long-term, basic research on future computing technologies with relevance to the DOE mission.

Applied Mathematics

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 and libraries, multiscale and multiphysics modeling, AI/ML, 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, scientific AI/ML, and for enabling greater capabilities for scientific discovery, design, and decision-support in complex systems.

Computer Science

The Computer Science research activity 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. In addition, and in partnership with the other SC programs and their scientific user facilities, the Computer Science activity supports basic research that addresses the need to seamlessly and intelligently integrate simulation, data analysis, and other tasks into coherent and usable workflows— from the edge of experiments, through simulation and AI, to data analytics and visualization.

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 and data infrastructure 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 a key role in developing and evolving 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, a recognized leader for the employment of high-performance computing (HPC) for scientific discovery. Established in 2001, SciDAC involves ASCR collaboration with the other SC programs and other DOE program offices with a goal to dramatically accelerate progress in scientific computing through deep collaborations between discipline scientists, applied mathematicians, and computer scientists. SciDAC does this by providing the intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to SC and DOE.

The Computational Partnerships activity also supports collaborations in the areas of data analysis and future computing. Collaboratory and data analysis projects 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 research and development of software to support a distributed data infrastructure and computing environment. The partnerships with BES, BER, FES, HEP, and NP enable development of new algorithms and applications targeted for future computing platforms, including quantum information systems.

Advanced Computing Research

This activity supports Research and Evaluation Prototypes (REP) and the Computational Sciences Graduate Fellowship (CSGF).

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. As the challenges in this regime are increasingly linked with advances in the research program, this activity is realigned from the Facilities subprogram to the Research subprogram beginning in FY 2021.

In addition, this activity partners with the NNSA to support the CSGF.

Activities and Explanation of Changes

(dollars in thousands)

FY 2020 Enacted		FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
Mathematical, Computational, and Computer Sciences Research	**$155,000**	**$264,028**	**+$109,028**
Applied Mathematics	$41,500	$53,728	+$12,228
Funding supports Applied Mathematics 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/ML.		The Request will expand support of core research efforts in algorithms, libraries and methods that underpin high-end scientific simulations, scientific AI/ML techniques, and methods that help scientists extract insights from massive scientific datasets with an emphasis on foundational capabilities in AI/ML.	The increase will support new efforts that develop the advanced mathematics and methods to enable an integrated computational and data infrastructure for scientific discovery. It also will support expanded investments to develop foundational capabilities in scientific AI/ML, with a focus on the requirements of SC's critical applications and scientific datasets.
Computer Science	$38,700	$49,605	+$10,905
Funding supports Computer Science 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 SC user facilities. The activity also expands funding for efforts in quantum networking.		The Request will continue support for core investments in software that improves the utility of HPC and advanced networks for science, including AI techniques, workflows, tools, data management, analytics and visualizations with strategic increases focused on critical tools, including AI, to enable an integrated computational and data infrastructure. Funding for this activity will also expand long-term efforts that explore and prepare for emerging technologies, such as quantum networking, specialized and heterogeneous hardware and accelerators, quantum and neuromorphic computing.	The increase will support new efforts that develop workflows, software, and edge technologies to enable an integrated computational and data infrastructure for scientific discovery. Also, it will support expanded investments to develop scientific AI/ML software and tools with a focus on the requirements of SC's critical applications and scientific datasets. The increase will include additional funding for basic research in Quantum networking.

(dollars in thousands)

FY 2020 Enacted		FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
Computational Partnerships	$69,142	$75,051	+$5,909
Funding supports the SciDAC institutes and partnerships awarded in FY 2017- FY 2018, this activity maintains efforts in QIS in partnership with the other SC programs, and efforts to bring the power of HPC to data intensive science.		The Request will continue support for the SciDAC Institutes, and ASCR will recompete partnerships with SC and DOE applications. Partnerships on scientific data and AI will be continued with new partners added. Building on these efforts, the Request will support the foundations of a new integrated computational and data infrastructure for science that will more effectively and efficiently address SC's data needs. A new partnership with NIH will leverage DOE infrastructure to address the data analytics needs of the connectome project and ensure that data is widely available for SC's AI development efforts to incorporate the results. The Request also includes support for a partnership with BES, HEP, and FES on microelectronics research.	The increase will support expanded efforts in AI data partnerships, QIS partnerships and Centers, and the recompetition of the SciDAC applications in partnership with other SC Programs. It also supports new efforts in data infrastructure and analytics, in microelectronics, and in partnerships with SC and interagency efforts, to maximize the public benefit and accelerate development of AI techniques for science.
Advanced Computing Research	$—	$76,007	+$76,007
Research and Evaluation Prototypes	*$—*	*$66,007*	*+$66,007*
In FY 2020, REP is funded within the High Performance Computing and Network Facilities subprogram.		The Request will continue to support quantum testbed efforts, with emphasis on partnerships with the new QIS centers, as well as small investments in cyber security and testbeds for microelectronics research. Building on basic research in quantum information networks, ASCR will support early stage research associated with the first steps to establishing a dedicated Quantum Network.	To address the need for strong connections between core research efforts and beyond Moore's Law prototypes, this activity is realigned from the Facilities subprogram to the Research subprogram in FY 2021. The increase will support early stage research for a Quantum internet and new testbed activities in support of SC efforts to revolutionize microelectronics.

(dollars in thousands)

FY 2020 Enacted	FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
Computational Sciences Workforce Programs *$—*	*$10,000*	*+$10,000*
In FY 2020, Computational Science Graduate Fellowship (CSGF) is funded within the High Performance Computing and Network Facilities subprogram.	The Request will provide support for the CSGF fellowship at $10,000,000, in partnership with NNSA. The goal of CSGF to increase availability of a trained workforce for exascale, AI, and beyond Moore's Law capabilities such as QIS.	To address the need for strong connections between core research efforts and beyond Moore's Law prototypes, this activity is shifted from the Facilities subprogram to the Research subprogram in FY 2021. The fellowship will emphasize high performance computational applications of AI and QIS.
SBIR/STTR $5,658	$9,637	+$3,979
In FY 2020, SBIR/STTR funding is set at 3.65 percent of non-capital funding.	In FY 2021, SBIR/STTR funding is set at 3.65 percent of non-capital funding.	The SBIR/STTR funding will be consistent with the ASCR total budget.

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), which provides high performance computing resources and large-scale storage to a broad range of SC researchers; and the two Leadership Computing Facilities (LCFs) at Oak Ridge National Laboratory (ORNL) and Argonne National Laboratory (ANL), which 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 include investments in upgrades: for the HPC user facilities, this scope includes electrical and mechanical system enhancements to ensure each remains state-of-the-art and can install future systems; for ESnet, the upgrades include rolling capacity growth to ensure no bottlenecks occur in the network. As the challenges within the Research and Evaluation Prototypes (REP) activity are increasingly linked with advances in the research program, this activity is being realigned to the Mathematical, Computational, and Computer Sciences Research subprogram beginning in FY 2021.

ASCR regularly gathers requirements from the other SC research programs through formal processes, including workshops and technical reviews, 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 and data trends, especially as the computing industry moves toward exascale computing and explores new architectures and technologies.

Allocation of computer time at ASCR HPC facilities follows the merit review public-access model used by other SC scientific user facilities. Two of ASCR's allocation programs include the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects for access to the LCFs and the ASCR Leadership Computing Challenge (ALCC) projects to provide an allocation path for critical DOE mission applications and to broaden the community of users.

High Performance Production Computing

This activity supports NERSC at LBNL to deliver high-end production computing services for the SC research community. Approximately 7,500 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 percent based in universities, 46 percent in DOE laboratories, and 5 percent 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 petaflops (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. NERSC will transition NERSC-9 to operations in FY 2021 while initiating preparations for a NERSC-10 upgrade that is needed to meet SC's growing computational needs.

Leadership Computing Facilities

The LCFs enable open scientific applications, including industry applications, to harness the potential of leadership computing to advance science and engineering. The success of this effort is built on the gains made in REP and ASCR research efforts. Another LCF strength is the staff, who operate and maintain the forefront computing resources and provide support to INCITE and 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, November 2018, June 2019, and November 2019. INCITE applications at Summit include: simulating how neutron star collisions produce heavy elements like gold and platinum; understanding how drug receptors select which signaling proteins to activate so as to enable the development of finely tuned medicines that yield desired effects with fewer side effects; closing, evaluating, and validating multiphase flow models in porous medium systems; new insights into the mechanisms leading to the complex phases and physical behavior observed in unconventional superconductors and quantum spin liquids and Monte Carlo simulations that will provide high-accuracy data for the adsorption of water on graphene with potential applications in water purification, desalinization and drug delivery. 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 will undertake final preparations for the deployment of a Cray/AMD exascale system to be named Frontier 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, to be known as Aurora. Aurora, will be the Nation's first exascale system when deployed in calendar year 2021 and is being designed by Intel/Cray to support the largest-scale computational simulations possible as well as large-scale data analytics and machine learning. INCITE applications at the ALCF include developing an understanding of the structure and reactions of nuclei to guide new experiments at the Facility for Rare Isotope Beams and at Thomas Jefferson National Accelerator Facility; identifying novel therapies to rationally design new treatments for a broad range of human cancers; discovering how water cycles affects river flow and freshwater supplies; and increasing the fidelity of earthquake models to improve the accuracy of seismic hazard assessment. 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 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. The demand for 2018 INCITE allocations at the LCFs outpaced the available resources by more than a factor of two.

Research and Evaluation Prototypes

To address the need for strong connections between core research efforts and beyond Moore's Law prototypes, this activity is shifted to the Mathematical, Computational, and Computer Sciences Research subprogram under the Advanced Computing Research activity in FY 2021.

High Performance Network Facilities and Testbeds

The Energy Sciences Network (ESnet) is SC's high performance network user facility, delivering highly reliable data transport capabilities optimized for the requirements of large-scale science. 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 that spans the continental United States and the Atlantic Ocean. ESnet interconnects all 17 DOE national laboratories, dozens of other DOE sites, and approximately 200 research and commercial networks around the world, enabling many 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 percent 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. ESnet is currently designing and executing a complete upgrade of its backbone network (the ESnet-6 upgrade).

Activities and Explanation of Changes

(dollars in thousands)

FY 2020 Enacted		FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
High Performance Computing and Network Facilities	**$636,265**	**$555,078**	**-$81,187**
High Performance Production Computing	$110,000	$85,000	-$25,000
Funding continues 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.		The Request will support operations at the NERSC facility, including user support, power, space, system leases, and staff. The Request will also support completion and transition to operations for the NERSC-9 upgrade, including site preparation activities, system acquisition, and application readiness.	Funding will support the operations of the Cori and NERSC-9 systems as well as early stage designs for the NERSC-10 upgrade.
Leadership Computing Facilities	$375,000	$370,000	-$5,000
Funding supports 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, are also supported.		The Request will support operations at the LCF facilities at ANL and ORNL, including user support, power, space, system leases, and staff. The Request also will support final site preparation for the ALCF-3 upgrade and OLCF-5 upgrade, and early access system testbeds.	Funding will support the 2021-2022 deployment schedule for both LCF upgrades. The decrease at the OLCF will result in the shifting of some milestone payments for Frontier to the lease agreement.
Argonne Leadership Computing Facility (ALCF)	*$150,000*	*$150,000*	*$—*
Funding supports continual operation of Theta and testbeds that were deployed in FY 2019 to support SC-ECP. 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 calendar year 2021 timeframe as part of ECI, while decommissioning MIRA.		The Request will continue support for the operation and competitive allocation of the Theta system. In support of ECP, the ALCF will provide access to Theta and other testbeds for ECP application and software projects. The ALCF will continue activities to enable deployment of the ALCF-3 exascale system, Aurora in the calendar year 2021 timeframe under CORAL I.	Funding will continue operations at the ALCF.

(dollars in thousands)

FY 2020 Enacted		FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
Oak Ridge Leadership Computing Facility (OLCF)	*$225,000*	*$220,000*	*-$5,000*
Funding supports continual operation and allocation of Summit. In support of ECP, the OLCF provides access to Summit and other testbeds 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 Request will continue support for the operation and competitive allocation of the Summit system. In support of ECP, the OLCF will provide access to Summit and other testbeds for ECP application and software projects. The OLCF will continue activities to enable deployment of the OLCF-5 exascale system, Frontier in the calendar year 2021-2022 timeframe under CORAL II.	The reduction at the OLCF will result in the shifting of milestone payments for Frontier to the lease agreement.
Research and Evaluation Prototypes	$39,000	$—	-$39,000
Funding supports 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.		To address the need for strong connections between core research efforts and beyond Moore's Law prototypes, this activity is realigned to the Advanced Computing Research activity within the research subprogram in FY 2021.	This reduction reflects the realignment of this program to the Research subprogram.
High Performance Network Facilities and Testbeds (ESnet)	$90,000	$80,000	-$10,000
Funding supports operations of the ESnet at 99.9 percent reliability. In addition, funding supports the ESnet-6 upgrade to increase network capacity and modernize the network architecture.		The Request will support operations of ESnet at 99.9 percent reliability, including user support, operations and maintenance of equipment, fiber leases, R&D testbed, and staff. The Request will continue support for the ESnet-6 upgrade to build the next generation network with new equipment, increased capacity, and an advanced network architecture.	Funding will support the upgrade in accordance with the project baseline as well as continued operation of ESnet-5.
SBIR/STTR	$22,265	$20,078	-$2,187
In FY 2020, SBIR/STTR funding is set at 3.65 percent of non-capital funding.		In FY 2021, SBIR/STTR funding is set at 3.65 percent of non-capital funding.	The SBIR/STTR funding will be consistent with the ASCR total budget.

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 deployment of these systems includes necessary site preparations and non-recurring engineering (NRE) at the Leadership Computing Facilities (LCF) that will ultimately house and operate the exascale systems.

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. 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 2021 Request includes $168,945,000 for the SC-ECP. These funds will support the final preparation of mission critical applications and development of a software stack for exascale platforms, and the initiation of testing on exascale early access platforms. Deployment of exascale systems in calendar years 2021-2022 will be through the LCFs as part of their usual upgrade processes.

Activities and Explanation of Changes

(dollars in thousands)

FY 2020 Enacted		FY 2021 Request	Explanation of Changes FY 2021 Request vs FY 2020 Enacted
Construction	**$188,735**[a]	**$168,945**	**-$19,790**
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	$188,735	$168,945	-$19,790
Funding supports the acceleration of application and software stack development in preparation for delivery of the first exascale system in calendar year 2021.		The Request will support project management; co-design activities between application and the software stack; and integration between SC-ECP and the LCF to provide continuous integration and testing of the ECP funded applications and software on exascale testbed.	Funding will decrease with the completion of vendor partnerships as well as the project's focus moving more to the final preparations for testing and final execution on the exascale platforms following acceptance rather than on exploratory research.

[a] In addition, $275,000,000 of ECI funding is provided within the Leadership Computing Facilities activity in FY 2020 and $270,000,000 is requested in FY 2021 in the Leadership Computing Facilities activity 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 2019 Enacted	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
Capital Operating Expenses						
Capital Equipment	N/A	N/A	5,000	5,000	5,000	—
Total, Capital Operating Expenses	**N/A**	**N/A**	**5,000**	**5,000**	**5,000**	**—**

Capital Equipment

	(dollars in thousands)					
	Total	Prior Years	FY 2019 Enacted	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
Capital Equipment						
Total, Non-MIE Capital Equipment	N/A	N/A	5,000	5,000	5,000	—
Total, Capital Equipment	**N/A**	**N/A**	**5,000**	**5,000**	**5,000**	**—**

Funding Summary

	(dollars in thousands)			
	FY 2019 Enacted	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
Research	408,500	405,000	453,051	+48,051
Facility Operations	527,000	575,000	535,000	-40,000
Total, Advanced Scientific Computing Research	**935,500**	**980,000**	**988,051**	**+8,051**

Advanced Scientific Computing Research
Scientific User Facility Operations

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

Definitions 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 2019 Enacted	FY 2019 Current	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
TYPE A FACILITIES					
NERSC	**$104,000**	**$104,000**	**$110,000**	**$85,000**	**-$25,000**
Number of users	7,500	7,500	7,500	7,500	—
Achieved operating hours	8,482	8,482	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 2019 Enacted	FY 2019 Current	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
OLCF	**$199,000**	**$199,000**	**$225,000**	**$220,000**	**-$5,000**
Number of users	1,450	1,450	1,450	1,450	—
Achieved operating hours	7,000	7,000	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
ALCF	**$140,000**	**$140,000**	**$150,000**	**$150,000**	**$—**
Number of users	950	950	950	950	—
Achieved operating hours	6,945	6,945	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	**$84,000**	**$84,000**	**$90,000**	**$80,000**	**-$10,000**
Number of users	N/A	N/A	N/A	N/A	—
Achieved operating hours	8,760	8,760	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[a]	**$527,000**	**$527,000**	**$575,000**	**$535,000**	**-$40,000**
Number of users	9,900	9,900	9,900	9,900	—
Achieved operating hours	31,187	31,187	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 optimal hours[b]	99%	99%	N/A	N/A	N/A
Unscheduled downtime hours	1%	1%	N/A	N/A	N/A

[a] ASCR prioritizes operations of the facilities. Annual budget shifts at ASCR facilities are driven by or impact upgrade projects and do not change operating hours.

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

Advanced Scientific Computing Research
Scientific Employment

	FY 2019 Enacted	FY 2020 Enacted	FY 2021 Request	FY 2021 Request vs FY 2020 Enacted
Number of permanent Ph.D.'s (FTEs)	634	683	783	+100
Number of postdoctoral associates (FTEs)	315	331	350	+19
Number of graduate students (FTEs)	407	438	525	+87
Other scientific employment (FTEs)[a]	196	212	205	-7

[a] 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

The FY 2017 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 $270,000,000 in FY 2021 to support the final site preparations and the start of installation of the exascale systems at both the Argonne and Oak Ridge Leadership Computing Facilities (LCFs). 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 costs of exascale systems, which is not included in the SC-ECP, will be funded within the ASCR facility budgets in the outyears. This Project Data Sheet (PDS) is for the SC-ECP only; prior-year activities related to the SC-ECP are also included.

The FY 2021 Request for SC-ECP will support project management; co-design activities between application development, software technologies, and hardware node and system design; continuous integration and testing of ECP-funded application and software tools and technologies onto SC LCF pre-exascale resources; and increased engagement and integration between SC-ECP and the LCFs in readying the ECP software stack and applications via testing and development on early access exascale hardware.

Significant Changes

The FY 2021 Request for SC-ECP is $168,945,000. The FY 2021 Request supports investments in the ECP technical focus areas—application development, software technology and hardware and integration—to finalize a capable exascale software ecosystem that supports the successful delivery of the first exascale-capable computer in the calendar year 2021 timeframe. The decrease reflects the completion of vendor partnerships as well as project focus moving more to development and deployment rather than on research.

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. The project is scheduled to achieve CD-2 in the second quarter of FY 2020. The estimated Total Project Cost (TPC) range of the SC-ECP is $1.0 billion to $2.7 billion.

Critical Milestone History

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	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	7/28/16	2Q FY 2019	1/03/17	4Q FY 2019	3Q FY 2019	4Q FY 2019	N/A	4Q FY 2023
FY 2019	7/28/16	2Q FY 2019	1/03/17	4Q FY 2019	3Q FY 2019	4Q FY 2019	N/A	4Q FY 2023

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	CD-2	Final Design Complete	CD-3B	D&D Complete	CD-4
FY 2020	7/28/16	2Q FY 2019	1/03/17	1Q FY 2020	3Q FY 2019	1Q FY 2020	N/A	4Q FY 2023
FY 2021	7/28/16	3/22/16	1/03/17	2Q FY 2020	6/06/19	2Q FY 2020	N/A	4Q FY 2024

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 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-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	CD-3A
FY 2017	TBD	TBD
FY 2018	4Q FY 2019	1/03/17
FY 2019	4Q FY 2019	1/03/17
FY 2020	1Q FY 2020	1/03/17
FY 2021	1Q FY 2020	1/03/17

CD-3A – Approve Long Lead Time Procurements

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
FY 2021	N/A	507,680	507,680	818,526	N/A	818,526	1,326,206

2. Project Scope and Justification

Scope

Four well-known challenges[a] are key to requirements and Mission Need 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://www.isgtw.org/feature/opinion-challenges-exascale-computing

The realization of an exascale-capable system that addresses parallelism, resilience, energy consumption, and memory/storage involves 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 U.S. HPC 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 the ECI.

Justification

The SC-ECP is 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 were focused on research in software (new algorithms and methods to support application and system software development) and hardware (node and system design), and these costs will be reported as Other Project Costs. During the last three years of the project, activities will focus primarily on hardening the application and the system stack software, and on additional hardware technology 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. The Objective KPPs represent the desired project performance. Achievement of the Threshold KPPs will be a prerequisite for approval of CD-4, Project Completion.

Performance Measure	Threshold	Objective
Exascale performance improvements for mission-critical challenge problems	50 percent of selected applications achieve Figure of Merit improvement greater than or equal to 50x	100 percent of selected applications achieve Figure of Merit improvement greater than or equal to 50x
Broaden exascale science and mission capability	50 percent of the selected applications can execute their challenge problem[a]	100 percent of selected applications can execute their challenge problem
Productive and sustainable software ecosystem	50 percent of the weighed impact goals are met	100 percent of the weighted impact goals are met
Enrich the HPC Hardware Ecosystem	Vendors meet 80 percent of all the PathForward milestones	Vendors meet 100 percent of all the PathForward milestones

[a] 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.

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Construction			
FY 2020	174,735	174,735	174,735
FY 2021	154,945	154,945	154,945
Outyears	178,000	178,000	178,000
Total, TEC	**507,680**	**507,680**	**507,680**
Other Project Cost (OPC)			
(Research for Application Development, System Software Technology, and Hardware Technology)			
FY 2016[a]	146,820	146,820	2,191
FY 2017	164,000	164,000	90,425
FY 2018	205,000	205,000	177,363
FY 2019	232,706	232,706	216,728
FY 2020	14,000	14,000	136,040
FY 2021	14,000	14,000	153,779
Outyears	42,000	42,000	42,000
Total, OPC	**818,526**	**818,526**	**818,526**
Total Project Cost (TPC)			
FY 2016[a]	146,820	146,820	2,191
FY 2017	164,000	164,000	90,425
FY 2018	205,000	205,000	177,363
FY 2019	232,706	232,706	216,728
FY 2020	188,735	188,735	310,775
FY 2021	168,945	168,945	308,724
Outyears	220,000	220,000	220,000
Total, TPC	**1,326,206**	**1,326,206**	**1,326,206**

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

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.

(dollars in thousands)

	Current Total Estimate[a]	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Application Development	236,759	211,800	N/A
Production Ready Software	142,661	147,190	N/A
Hardware Partnership	128,260	67,745	N/A
Total, TEC	**507,680**	**426,735**	**N/A**
Other Project Cost (OPC)			
Planning/Project Mgmt	102,595	123,715	N/A
Application Development	333,568	298,824	N/A
Software Research	199,754	167,970	N/A
Hardware Research	182,609	239,141	N/A
Total, OPC	**818,526**	**829,650**	**N/A**
Total Project Cost	**1,326,206**	**1,256,385**	**N/A**

5. Schedule of Appropriations Requests

(dollars in thousands)

Request Year	Type	Prior Years[b]	FY 2019	FY 2020	FY 2021	Outyears	Total
FY 2017	TEC	—	TBD	TBD	TBD	TBD	TBD
	OPC	311,894	TBD	TBD	TBD	TBD	TBD
	TPC	311,894	TBD	TBD	TBD	TBD	TBD
FY 2018	TEC	—	—	175,000	—	215,000	390,000
	OPC	518,524	189,000	14,000	—	42,000	763,524
	TPC	518,524	189,000	189,000	—	257,000	1,153,524
FY 2019	TEC	—	—	174,735	—	252,000	426,735
	OPC	518,524	232,706	14,000	—	42,000	807,230
	TPC	518,524	232,706	188,735	—	294,000	1,233,965
FY 2020	TEC	—	—	174,735	—	252,000	426,735
	OPC	526,944	232,706	14,000	—	56,000	829,650
	TPC	526,944	232,706	188,735	—	308,000	1,256,385
FY 2021	TEC	—	—	174,735	154,945	178,000	507,680
	OPC	515,820	232,706	14,000	14,000	42,000	818,526
	TPC	515,820	232,706	188,735	168,945	220,000	1,326,206

[a] Estimate includes distribution of contingency funds based on percentage of planned cost; these values will be adjusted as contingency is distributed.

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

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 Baseline Change Proposal (BCP) was executed in March 2018 to reflect this change. In the FY 2021 Budget Request, $270,000,000 is included in the Argonne and Oak Ridge National Laboratories' LCF 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 or Beneficial Occupancy	FY 2022
Expected Useful Life	5 years
Expected Future Start of D&D of this capital asset	4Q 2030

7. D&D Information

N/A, no construction.

8. Acquisition Approach

The early years of the SC-ECP, approximately four years in duration, supported 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.