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 innovative 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 deploys and operates world-class, open access high performance computing facilities and a high performance network infrastructure 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 advanced computing systems and 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. 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 innovative clean energy technologies, medicines, industries and supply chains, and military capabilities. The U.S. will need to leverage investments in science for innovative new technologies, materials, and methods to build back better. Most of the modeling and prediction necessary to produce the next generation of breakthroughs in science will come from employing data-driven methods at extreme scale tightly coupled to the enormous increases in the volume and complexity of data generated by U.S. researchers and SC user facilities. The convergence of AI technologies with these existing investments creates a powerful accelerator for innovation.

The emerging field of 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. 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. However, 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.

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.

ASCR's proposed activities will advance AI, QIS, advanced communication networks, and strategic computing to accelerate progress in delivering a clean energy future, understanding and addressing climate change and increasing the competitive advantage of U.S. industry.

Highlights of the FY 2022 Request

The FY 2022 Request of $1,040.0 million for ASCR will strengthen U.S. leadership in strategic computing with operation of the Nation's first exascale computing system and deployment of a second system, broadening the foundations of AI and QIS, and expanding the infrastructure that enables data-driven science from climate to clean energy solutions.

Research

- To ensure ASCR is meeting SC's HPC and advanced networking mission needs during and after the exascale project, the Request prioritizes foundational research in Applied Mathematics and Computer Science. Investments will continue to emphasize the challenges of data intensive science, including AI/ML, and development of innovative computing and networking technologies. The Request increases support for ASCR's Computational Partnerships with a focus on developing partnerships that broaden the impact of both exascale and data infrastructure investments in areas of strategic importance to DOE and the Nation leading to new clean energy solutions, greater understanding of the earth's systems, new accelerator technologies, and enhancing our ability to respond to national emergencies. This includes partnering with other programs and key agencies to understand their simulation and modeling capabilities, data management and curation needs, and to identify and bridge gaps to enable DOE to provide appropriate resources. The Request also increases support for the Computational Sciences Graduate Fellowship (CSGF) to increase the number of fellows in AI and Quantum as well as outreach to and participation by under-represented groups.

- The Request provides robust support for Advanced Computing Research's quantum investments in the National Quantum Information Sciences (QIS) Research Centers, quantum internet and testbeds. ASCR will continue to partner with the other SC programs to support the multi-disciplinary National QIS Research 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 2022, ASCR will support early-stage research and development associated with the first steps to establishing a dedicated Quantum Network.

- The Office of Science is fully committed to advancing a diverse, equitable, and inclusive research community. This commitment is key to providing the scientific and technical expertise for U.S. leadership in high end computing, networking, and computational science. Toward that goal, ASCR will participate in the SC-wide Reaching a New Energy Sciences Workforce (RENEW) initiative that leverages SC's unique national laboratories, user facilities, and other research infrastructures to provide undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem. This includes Historically Black Colleges and Universities (HBCUs) and Minority Serving Institutions (MSIs) and individuals from groups historically underrepresented in STEM, but also includes students from communities with environmental justice impacts and the EPSCoR jurisdictions. The hands-on experiences gained through the RENEW initiative will open new career avenues for the participants, forming a nucleus for a future pool of talented young scientists, engineers, and technicians with the critical skills and expertise needed for the full breadth of SC research activities, including DOE national laboratory staffing.

Facility Operations

- 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. Funding supports the baseline schedules for planned 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 at the facilities and provides robust support for the execution of ECI, which includes the SC-Exascale Computing Project (SC-ECP).

- Current ASCR high performance computing (HPC) resources and facilities are designed to efficiently execute large-scale simulations and are focused on minimizing users' wait-times in batch queues while maximizing use of these unique resources. However, the rate and volume of data from SC scientific user facilities is expected to grow exponentially in the future. ASCR will design a state-of-the-art Scientific HPC Data Facility focused on the unique challenges of near real-time computing needed to support the explosion of scientific data that will serve as the anchor for the integrated computational and data infrastructure efforts. To provide geographic diversity, this facility will be located on the East Coast.

Projects

- The ASCR FY 2022 Request includes $404.0 million for SC's contribution to DOE's Exascale Computing Initiative to deploy an exascale computing software ecosystem and mission critical applications on at least one exascale system in calendar year 2021 and a second in the 2021–2022 timeframe to address national needs. $275 million of this effort will go to the LCFs to deploy and operate the exascale systems and testbeds and support ECP project teams.

Advanced Scientific Computing Research
FY 2022 Research Initiatives

Advanced Scientific Computing Research supports the following FY 2022 Research Initiatives.

(dollars in thousands)

	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Artificial Intelligence and Machine Learning	36,000	56,866	58,820	+1,954
Biopreparedness Research Virtual Environment (BRaVE)	–	–	5,183	+5,183
Exascale Computing Crosscut	463,735	438,945	404,000	-34,945
Integrated Computational & Data Infrastructure	–	11,974	15,328	+3,354
Microelectronics	–	5,182	5,183	+1
Quantum Information Science	54,680	98,402	107,649	+9,247
Reaching a New Energy Sciences Workforce (RENEW)	–	–	5,000	+5,000
Total, Research Initiatives	**554,415**	**611,369**	**601,163**	**-10,206**

Advanced Scientific Computing Research
Funding

(dollars in thousands)

	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Advanced Scientific Computing Research				
Applied Mathematics Research	41,500	48,570	51,048	+2,478
Computer Sciences Research	38,700	46,827	49,773	+2,946
Computational Partnerships	69,142	76,194	86,029	+9,835
Advanced Computing Research	–	88,274	106,112	+17,838
Mathematical, Computational, and Computer Sciences Research, SBIR/STTR	5,658	–	–	–
Total, Mathematical, Computational, and Computer Sciences Research	**155,000**	**259,865**	**292,962**	**+33,097**
High Performance Production Computing	110,000	113,786	115,963	+2,177
Leadership Computing Facilities	375,000	381,075	408,113	+27,038
Research and Evaluation Prototypes	39,000	–	–	–
High Performance Network Facilities and Testbeds	90,000	91,329	93,962	+2,633
High Performance Computing and Network Facilities, SBIR/STTR	22,265	–	–	–
Total, High Performance Computing and Network Facilities	**636,265**	**586,190**	**618,038**	**+31,848**
17-SC-20 SC Exascale Computing Project	**188,735**	**168,945**	**129,000**	**-39,945**
Subtotal, Advanced Scientific Computing Research	**980,000**	**1,015,000**	**1,040,000**	**+25,000**
Total, Advanced Scientific Computing Research	**980,000**	**1,015,000**	**1,040,000**	**+25,000**

SBIR/STTR funding:

- FY 2020 Enacted: SBIR $25,160,000 and STTR $3,538,000
- FY 2021 Enacted: SBIR $25,736,000 and STTR $3,620,000
- FY 2022 Request: SBIR $30,753,000 and STTR $4,324,000

Advanced Scientific Computing Research
Explanation of Major Changes

	(dollars in thousands) FY 2022 Request vs FY 2021 Enacted
Mathematical, Computational, and Computer Sciences Research	**+$33,097**

The Computer Science and Applied Mathematics activities will continue to increase their efforts on foundational research to address 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 quantum, AI and data infrastructure tools, into strategic applications, including areas such as revolutionizing microelectronics, accelerating the development of clean energy technologies, and understanding the earth's systems, to get the most out of the leadership computing systems and data infrastructure investments. In addition, the Computational Partnerships activity will increase investments in the development of algorithms, applications, and data infrastructure, focused on both AI 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 and new efforts with other agencies to enable rapid response to deliver research capabilities to address national emergencies. Computer Science for quantum information networks will continue to focus on addressing new opportunities and challenges of transporting and storing quantum information. The Advanced Computing Research activity continues to robustly support the National QIS Research Centers and quantum testbeds, in close coordination with the other SC programs. This includes early-stage R&D activities for a quantum network. The Research and Evaluation Prototype activity also continues to support fundamental research in cybersecurity, microelectronics and emerging technologies such as neuromorphic computing. This subprogram also supports the RENEW initiative to provide undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem. In addition, the Computational Sciences Graduate Fellowship is increased to expand the skilled workforce necessary to maintain U.S. Leadership in high end computing, networking, and computational science.

High Performance Computing and Network Facilities	**+$31,848**

The LCFs will complete acceptance testing and early science/ECP access to the Frontier exascale computing system at Oak Ridge and the Aurora exascale system at Argonne through calendar year 2022. Both facilities will also provide testbed resources to the SC-ECP to continuously test and deploy software technologies. In addition, funding supports operation of the 125 petaflop NERSC-9 Perlmutter system, planning for NERSC-10, and the ESnet-6 upgrade to significantly increase capacity and security at all DOE sites. Funding for all facilities supports operations, including power, equipment, staffing, testbeds, and lease payments. To address the significant growth in the rate and volume of data from SC scientific user facilities, ASCR will design a state-of-the-art Scientific HPC Data Facility focused on the unique challenges of near real-time computing needed to support the explosion of data that will also serve as the anchor for the integrated computational and data infrastructure efforts.

	(dollars in thousands) FY 2022 Request vs FY 2021 Enacted
Exascale Computing The FY 2022 Request will support efforts to deploy SC-ECP applications and ecosystem on both exascale architectures in partnership with the ASCR facilities. The decrease represents transition from research to testing of applications and software on the exascale systems and testbeds.	**-$39,945**
Total, Advanced Scientific Computing Research	**+$25,000**

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 in AI, Future Advanced Computing 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 2016, 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 2022, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Supercomputing Versus COVID-19

In the early days of the pandemic it was clear that fighting COVID-19 would require extensive research in areas like bioinformatics, epidemiology, and molecular modeling to understand the threat and to develop strategies to address it. One of the earliest heavy hitters to join the research effort, OLCF's Summit supercomputer—at the time, the most powerful publicly ranked supercomputer in the world—quickly added projects to characterize the virus structure, understand the ways in which it infects cells and assess thousands of compounds for their ability to bind to SARS-CoV-2's spike protein. As allocation requests poured in and private and public computing resources from around the world offered help, the Department of Energy, with 16 founding partners, established the COVID-19 High Performance Computing Consortium in March 2020. This massive collaboration between government, academic, and industry partners from around the globe aimed to pool available resources—from small clusters to clouds and some of the worlds biggest supercomputers—and quickly match them to COVID researchers. IBM established a central portal for researchers to request resources and the consortium matched them with available computing resources from one of the partner institutions. An expert panel comprised of top scientists and computing researchers worked with proposers to assess the public health benefit of the work, with emphasis on projects that could ensure rapid results. The Consortium grew to 43 members with a computing capacity over 600 petaflops and received more than 175 research proposals from researchers in more than 15 countries. Projects used simulations to understand how the virus infects and the dynamics of COVID aerosols in various settings. They also used artificial intelligence to characterize the spread of the virus and identify effective countermeasures and to accelerate the discovery of promising treatments. The HPC consortium also worked with policymakers to manage the course of the infection and strategically deploy resources. The HPC consortium demonstrated how public-private partnerships can be quickly established to harness combined capabilities to address a global emergency and was recognized by HPC wire as the 2020 Editor's choice for "Best HPC Collaboration".

AI Joins the Fight Against COVID-19

To find a drug that can stop the SARS-CoV-2 virus, scientists want to screen billions of molecules for the right combination of properties. The process is usually risky and slow, often taking several years. However, a team of scientists led by Argonne National Laboratory found a way to make the process 50,000 times faster using artificial intelligence (AI). A related Argonne collaboration used a novel AI-driven workflow to elucidate the mechanism by which COVID infects cells in the paper, "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics", which was awarded the 2020 Special Gordon Bell prize for COVID applications at Supercomputing 2020. With that understanding and an expanded team, Argonne developed a pipeline of AI and simulation techniques to hasten the discovery of promising drug candidates for COVID-19. The pipeline is named IMPECCABLE, short for Integrated Modeling PipelinE for COVID Cure by Assessing Better Leads. With it, the team has been able to screen four billion potential drug candidates in a day—a thousand times more compounds that with conventional methods. This effort was part of the Department's National Virtual Biotechnology Laboratory, funded by the CARES Act, and leveraged the capabilities of Argonne's Advanced Photon Source (APS), the CARES Act funded enhancements to the Leadership Computing Facilities, the COVID-19 HPC Consortium, and capabilities developed within the DOE-NCI collaboration and the Exascale Computing Project.

Exascale Computing Project: Advancing the Applications for Building Back Better
The Exascale Computing Project (ECP) has been accelerating progress in a wide array of hardware and software technologies to advance National goals from basic science to clean energy options. Within the project, ECP researchers work with the ASCR facilities and domain scientists in the application development teams from across the Department's national labs and missions. ECP supports six application areas (national security, energy security, economic security, earth systems, and health care) and 24 applications with a significant focus on simulation and data-driven (AI) approaches. These efforts involve more than 10 billion lines of existing code that span methodologies and engage large user communities with hundreds of thousands of scientists and engineers standing to benefit from these efforts. Working together, in large diverse teams, they prepare and optimize critical applications for exascale architectures through knowledge of the application domain as well as the hardware, software, and other features of the planned exascale systems. All projects have ported to the OLCF's Summit architecture and have performance increases between a factor of 20 and a factor of 300 from ECP software improvements. The software advances, coupled with the exascale systems deployed in 2021–22 will enable these applications to deliver accurate regional impact assessment from climate change, forecast water resources and severe weather with increased confidence, address food supply changes, optimize power grid planning with diverse green energy options, reliably guide safe long-term consequential decisions about waste storage and carbon sequestration, accelerate the widespread adoption of additive manufacturing by enabling routine fabrication of qualifiable metal parts, design more robust and selective biofuel catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower, and enable atomistic simulations to assist in the development of novel materials for energy applications.

Exascale Computing Project: Partnering to Push U.S. Microelectronics Forward
For 50 years, steady progress in computing was described by "Moore's Law"––the observation that the number of transistors in an integrated circuit (IC) doubles about every two years. As feature sizes approached the width of individual atoms, the pace of progress was expected to slow. To sustain progress in scientific computing, ASCR initiated investments in joint research with U.S. chip makers to forge a new path forward. These "Forward" partnerships have, for example, enabled Advanced Micro Devices, Inc. (AMD) to provide immense computational power in America's pre-exascale and exascale supercomputers through innovative computing solutions, acceleration of key hardware and software technologies, the rejuvenation of the HPC hardware and software ecosystem, and joint research and co-design with the DOE National Laboratories. ECP partnered with AMD for collaborative research to accelerate and enhance the performance, power-efficiency, and capabilities of AMD's commercially available hardware and open-source software for high-performance computing and machine learning. One effort accelerated AMD's development and commercial use of its revolutionary chiplet architecture with an enhanced version of AMD's Infinity Fabric™ interconnect to link separate pieces of silicon, chiplets, within a single processor package. The use of chiplets provides higher performance and more cost-effective manufacturing than traditional monolithic chip designs, and enabled AMD to introduce the world's first high-performance x86 7nm CPUs. Chiplets and other technologies accelerated by the ECP "Forward" investments will be used to enable scientific breakthroughs in the Perlmutter pre-exascale supercomputer at the National Energy Research Scientific Computing Center (NERSC), the Frontier exascale supercomputer at the Oak Ridge Leadership Computing Facility, the NNSA's El Capitan exascale supercomputer at Lawrence Livermore National Laboratory (LLNL), and other supercomputers and datacenters across the Nation.

Exascale Computing Project: Accelerating the Future of U.S. Microelectronics
The demands of massive scientific datasets require ever increasing amounts of bandwidth to, within, and on High Performance Computers. Exascale systems, combined with the data from upgrades at many of the Office of Science scientific user facilities, push traditional interconnect technologies to their limits. Silicon Photonics, which uses photons of light to transmit information more rapidly and efficiently, offers the potential to be a disruptive technology for HPC, delivering significant performance gains similar to those realized by fiber optics in long-distance networks. This technology will also have significant impacts across microelectronics applications, leap-frogging foreign competition and helping to satisfy the increasing bandwidth demands of ubiquitous, online connectivity, including from 5G/6G deployment, and could even help deliver the promise of quantum computing, which would provide an additional disruptive technology. To address the challenges of data bottlenecks in exascale systems, ECP partnered with Intel to accelerate research in integrated Silicon Photonics. Intel improved the power efficiency of high-speed interconnects by developing co-packaged optical interconnect technology that integrates high-bandwidth Intel Silicon Photonic engines with next-generation high-bandwidth switches. The ECP-Intel "PathForward" project helped fund an intensive technology development effort that resulted in the demonstration of a fully functional Ethernet switch with co-packaged optics, demonstrating 400Gbps optical I/O links

interoperable with a commercial switching system. This demonstrates the technical viability of fully-integrated optical photonics with core compute silicon, and has motivated the formation of an Optical Interconnect Forum Co-Packaged Optical Working Group based upon the initial description published by leading data center partners (http://www.copackagedoptics.com/). The Working Group will publish industry-wide specifications for co-packaged optics by the end of 2021 to support a broad multi-vendor eco-system enabling broad technology deployment. The results of these efforts which will soon put the technology into HPCs, data centers, and telecommunication infrastructure that drive our modern information economy with performance and energy gains that are well beyond next-generation electrical-networking technologies.

Quantum Supremacy Milestone Harnesses OLCF Summit Supercomputer
While still in the early stages, quantum computers have the potential to be exponentially more powerful than today's leading high-performance computing systems with the promise to revolutionize research in materials, chemistry, high-energy physics, and across the science and technology spectrum. Quantum computers use the laws of quantum mechanics to increase greatly how information is processed. Demonstrating when this threshold in relative computational power is crossed has been an important event for the computing community known as quantum supremacy. A joint research team from Google Inc., NASA Ames Research Center, and Oak Ridge National Laboratory (ORNL) has demonstrated quantum supremacy using a Google quantum computer, Sycamore, which can outperform the world-class Summit supercomputer at the task of random circuit sampling (RCS). The RCS task was designed specifically to measure the performance of quantum devices by characterizing the rate at which a randomly selected distribution can be calculated. The Summit supercomputer was used to verify the accuracy and operation of the Google quantum computer against baseline RCS tasks. Then, running on more difficult case studies, Sycamore was shown to require 200 seconds to perform RCS whereas the same simulations on Summit was expected to require 10,000 years to complete. In addition, the calculation on Sycamore was found to be approximately 10 million times more energy efficient. These dramatic differences in performance have provided the first experimental evidence of quantum supremacy and give critical insights into the design of future quantum computers that can address scientific computing tasks inaccessible today.

The Power of Technology Convergence: HPC + AI + Physical Model = Gordan Bell Prize
Molecular dynamics modeling has become a primary tool in scientific inquiry, allowing scientists to analyze the movements of interacting atoms over a set period of time, which helps them determine the properties of different materials or organisms. These computer simulations often lead the way in designing everything from new drugs to improved alloys. However, the two most popular methodologies have tradeoffs between speed and accuracy and are limited by even the most powerful supercomputer. But what if there was a way to bridge the gap between these methods to produce complex simulations that are both large and accurate? With the power of OLCF's Summit supercomputer, researchers successfully tested a software package that offers a potential solution: DeePMD-kit, named for "deep potential molecular dynamics." The team refers to DeePMD-kit as a "HPC+AI+Physical Model" in that it combines high-performance computing (HPC), Artificial Intelligence (AI), and physical principles to achieve large-scale speed and accuracy. It uses a neural network to assist its calculations by approximating the physics, thereby reducing the computational complexity. Simulating a block of copper atoms, the team put DeePMD-kit to the test on Summit with the goal of seeing how far they could push the simulation's size and timescales. The team was able to simulate a system of 127.4 million atoms—more than 100 times larger than the current state of the art. Furthermore, the simulation achieved a time-to-solution mark of at least 1,000 times faster at 2.5 nanoseconds per day for mixed-half precision, with a peak performance of 275 petaflops (one thousand million million floating-point operations per second) for mixed-half precision. This result was awarded the 2020 Gordon Bell prize for outstanding achievements in high-performance computing at the 2020 Supercomputing Conference, SC20, and their methods will doubtless accelerate simulations in a wide array of applications.

DOE and NCI: Working Together Against Cancer
The mechanism and dynamics of how RAS proteins—a family of proteins whose mutations are linked to more than 30 percent of all human cancers—interact and promote cancer signaling are not well understood. A pilot project in the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) program, a collaboration between the Department of Energy (DOE) and National Cancer Institute (NCI) that is supported in part by the Cancer Moonshot™, was recognized with the Best Paper award at SC19 for a novel application of Artificial Intelligence to improve multi-scale Molecular Dynamics modeling. The paper describes a predictive approach to model the dynamics of RAS proteins and lipid membranes, as well

as the activation of oncogenic signaling through interaction with other proteins. The paper includes contributions from Lawrence Livermore National Laboratory (LLNL), Oak Ridge National Laboratory (ORNL), the Frederick National Laboratory for Cancer Research (FNLCR), and IBM. The team took a broad approach to modeling RAS protein interactions that validates the results of many smaller experiments and pushes toward a first of its kind model that fully describes all possible RAS interactions. They began with a macro-model capable of simulating the impact of a lipid membrane on RAS proteins at long timescales and incorporated a machine learning algorithm to determine which lipid "patches" were interesting enough to model in more detail with a molecular level micromodel. The result is a Massively parallel Multiscale Machine-Learned Modeling Infrastructure (MuMMI) that scales up efficiently on supercomputers like the OLCF's Summit. The models help NCI carry out experiments to test predictions and generate more data that will feed back into the machine learning model, creating a validation loop that will produce a more accurate model. The MuMMI tool is also being incorporated into DOE multiscale codes to improve predictive capability in energy and national security applications.

Machine Learning Accelerates Our Clean Energy Future
Magnetically confined fusion promises to deliver CO_2-free energy. Test runs on ITER is the next step in the development of fusion energy. The ITER walls surrounding the fusion plasma are made of tungsten because it can withstand the very high temperature of the plasma. However, the tungsten can also contaminate and subsequently degrade the plasma. Attempts to simulate the tungsten-plasma interactions, including the 74 different charge states in the ITER plasma, would require an entire exascale computing system. A new method, developed though a SciDAC collaboration between fusion physicists and applied mathematicians, uses machine learning to simulate the contamination process. This approach has been demonstrated to achieve the desired accuracy while significantly reducing the computing requirements from nearly all of an exascale computer to a few percent.

Toward Environmental Justice: How Urbanization and Pollution Increase Storm Activity Around Cities
A study led by scientists at Pacific Northwest National Laboratory (PNNL) used computing resources at the National Energy Research Scientific Computing Center (NERSC) to run detailed simulations of storm physics and found that urbanization, combined with human-produced (anthropogenic) aerosol air pollution, increased both storm strength and storm development in two major U.S. cities. The researchers used NERSC's Cori system to study two separate convective storms: a supercell near Kansas City, Missouri that produced hail, tornado, and strong wind; and a sea-breeze-induced thunderstorm occurring near Houston, Texas. To resolve convection patterns, a key part of the study, the simulations were run at very high resolution across the entire Cori supercomputer. With this resolution, researchers were able to gain a good understanding of how physical factors affect storms. They found that urban land modifies convective evolution, speeding up cloud state transitions and initiating surface rain earlier, all of which results from urban heating-induced stronger sea-breeze circulation. The anthropogenic aerosol effect became evident after the storm clouds evolved, accelerating convective intensity and precipitation mainly by activating numerous ultrafine particles at various cloud stages. The team is building on this understanding to build a more predictive modeling, with the goal of being able to foresee hazardous weather and prevent death and damages from severe storms.

Deep Learning Nurtures the Roots of Innovations
When engines fail, technicians struggle to know why. Understanding the root cause of scientific phenomena, or an industrial failure event, could significantly accelerate discoveries, save on maintenance costs, and drive innovative instrument/scientific simulation design. However, there are often long, complex, and nonlinear interactions that occur between the root causes and its effects, which makes it difficult to identify a true root cause. The ASCR-funded extreme-scale spatio-temporal learning team led by researchers at Brookhaven National Laboratory (BNL) formulated an approach to sort and understand complex interactions via high performance computing modeling and simulation of event relationships and interactions in a bidirectional (forward and backward) deep learning framework. The approach was evaluated on industrial applications, such as engine failure sensor data, and can be applicable to any space and time correlated datasets and experimental sensor networks.

Connecting the dots – From Instruments to HPC to Publications: Getting More Out of Our Investments in Science
Exponential increases in data volumes and velocities are overwhelming finite human capabilities. Continued progress in science and engineering demands that we automate a broad spectrum of currently manual research data manipulation tasks, from data transfer and sharing to acquisition, publication, and analysis. These needs are particularly evident in large-scale experimental science, in which researchers are typically granted short periods of instrument time and must maximize

experiment efficiency as well as output data quality and accuracy. To address the need for improved and automated workflows, the Robust Analytical Models for Extreme-Scale Systems (RAMSES) project is developing end-to-end analytical performance modeling that is transforming science workflows in extreme-scale science environments. In one example that was recognized at SC19, the team automated a real-world scientific workflow in which the Advanced Photon Source was used to image centimeter-scale mouse brains with sub-micrometer resolution acquiring real-time data that was automatically passed through a complex automation flow for reconstruction using HPC resources, human-in-the-loop coordination, and data publication and visualization. The team is building on lessons learned from that experiment to productionize the platform so it can be applied to a variety of use cases and scaled to a range of compute resources matched to the size of the datasets up to exascale computing for challenges such as the mouse brain connectome.

ESnet: Advancing the Democratization of Science
Over the last decade, the scientific community has experienced an unprecedented shift in the way research is performed and how discoveries are made. Highly sophisticated experimental instruments are creating massive datasets for diverse scientific communities that hold the potential for new insights that will have long-lasting impacts on society. However, scientists cannot make effective use of this data if they are unable to move, store, and analyze it. The Engagement and Performance Operations Center (EPOC) was established in 2018 as a collaborative focal point for operational expertise and analysis and is jointly led by Indiana University (IU) and the Energy Sciences Network (ESnet) with support from the National Science Foundation. EPOC provides researchers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. By considering the full end-to-end data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process. Many researchers at larger educational institutions, or part of large-scale collaborations, already have access to significant in-house resources, so EPOC focuses on small or medium-sized institutions and collaborations that may lack the financial and human resource capacity for more advanced services. By working with regional networks to develop, teach, and make available additional instructive material to these institutions, EPOC is not only increasing the abilities of the teams they are in direct contact with, but is also providing a broad set of materials made freely available to the general public.

Advancing Research in Geothermal Resources Mapping Using Distributed Acoustic Sensing
At present, large portions of western basins relevant to geothermal energy production are poorly mapped using classical high-resolution geophysical methods due to the high costs of seismic surveys, the restricted availability of archival seismic lines in the same regions, and limited coverage by dense arrays required for ambient noise imaging. ESnet is collaborating with the scientists to explore the use of fiber optic sensing, particularly distributed acoustic sensing (DAS), to seismically characterize geothermal systems at the basin scale in California’s Imperial Valley south of the Salton Sea, near Calipatria and El Centro. ESnet has leveraged its extensive experience in the telecom fiber marketplace to acquire access to specific dark fiber segments motivated by the research goals and rack space and power to house the DAS equipment. These efforts have led to a recent award under the DOE’s Frontier Observatory for Research in Geothermal Energy (FORGE) initiative.

ESnet: Advancing Diversity in Network Engineering
2020 marked the five-year anniversary of the Women in IT Networking at SC (WINS) program. ESnet created the program with collaborators funded by the NSF, to address the gender diversity gap in the network engineering workforce. The program funds and mentors approximately 7 women each year to participate in the construction of SCinet, one of the world’s most advanced networks, purpose-built for the annual Supercomputing Conference. SCinet provides these women the opportunity to work with ~150 leading engineers and top technology vendors to build a Terabit+ network. The awardees are selected through a competitive application process. Since its introduction, WINS has increased the female to male participation ratio in SCinet from 13 percent to over 30 percent as of 2020. In recent years, the program has funded women not just from large universities and DOE laboratories but also from community and tribal colleges.

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. ASCR investments are not focused on the next quarter but on the next quarter century. The Computer Science and Applied Mathematics activities in ASCR provide the foundation for increasing the capability of the national HPC ecosystem and scientific 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 the Nation. 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 incorporating 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 all levels of computing, from the edge to HPC storage and networking resources;
- collaboration tools, data and compute infrastructure and partnerships to make scientific resources and data 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, DOE, and the Nation; and
- long-term, basic research on future computing technologies with relevance to the DOE missions.

Applied Mathematics Research

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 multi-physics 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. Historically, advances in these methods have contributed as much, if not more, to gains in computational science than hardware improvements alone. Forward-looking efforts by this activity anticipate DOE mission needs from the closer coupling and integration of scientific modeling, data and scientific AI/ML with advanced computing, for enabling greater capabilities for scientific discovery, design, and decision-support in complex systems and new algorithms to support data analysis at the edge of experiments and instruments.

Computer Sciences Research

The Computer Science research activity supports long-term, basic research on the software infrastructure that is essential for the effective use of the most powerful high performance computing and networking systems in the country as well as the tools and data infrastructure to enable the exploration and understanding of extreme scale and 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 Computer Science research plays a key role in developing and evolving the sophisticated software required for future Leadership Computers, including basic research focused on quantum computing and communication. Hardware and software vendors often take software developed with ASCR Computer Science investments and integrate it with their own software. ASCR-supported activities are entering a new paradigm driven by sharp increases in the heterogeneity and complexity of computing systems and their software ecosystems, support for large-scale data analytics, and by the incorporation of AI techniques. 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 comprehensive workflows—from the edge of experiments, through simulation and AI, to data analytics and visualization. This includes making research data and AI models findable, accessible, interoperable, and reusable to strengthen trust and maximize the impact of scientific research in society.

Computational Partnerships

The Computational Partnerships activity primarily supports the Scientific Discovery through Advanced Computing, or SciDAC, program, which is a recognized leader for the employment of HPC for scientific discovery. Established in 2001, SciDAC involves ASCR partnerships with the other SC programs, other DOE program offices, and other federal agencies in strategic areas 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, DOE, and the Nation.

The Computational Partnerships activity also supports collaborations in the areas of data analysis, microelectronics, and future computing. Collaborative 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. In addition, 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 Quantum efforts, Research and Evaluation Prototypes (REP), and ASCR-specific workforce investments including the Computational Sciences Graduate Fellowship (CSGF) and the SC-wide Reaching a New Energy Sciences Workforce (RENEW) initiative.

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. In addition to REP, this activity supports ASCR’s investments in the National QIS Research Centers, as well as quantum computing testbeds and building a quantum internet to connect the National QIS Research Centers and ultimately the 17 DOE National Laboratories.

SC is fully committed to advancing a diverse, equitable, and inclusive research community, key to providing the scientific and technical expertise for U.S. scientific leadership. Toward that goal, ASCR will participate in the SC-wide RENEW initiative that leverages SC’s world-unique national laboratories, user facilities, and other research infrastructures to provide undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem. This includes HBCUs and MSIs and individuals from groups historically underrepresented in STEM, but also includes students from communities with environmental justice impacts and the EPSCoR jurisdictions. The hands-on experiences gained through the RENEW initiative will open new career avenues for the participants, forming a nucleus for a future pool of talented young scientists, engineers, and technicians with the critical skills and expertise needed for the full breadth of SC research activities, including DOE national laboratory staffing.

Success in fostering and stewarding a highly skilled, diverse, equitable, and inclusive workforce is fundamental to SC's mission and key to also sustaining U.S. leadership in HPC and computational science. The high demand across DOE missions and the unique challenges of high-performance computational science and engineering led to the establishment of the CSGF in 1991. This program has delivered leaders in computational science both within the DOE national laboratories and across the private sector. With increasing demand for these highly skilled scientist and engineers, ASCR continues to partner with the NNSA to support the CSGF to increase the availability and diversity of a trained workforce for exascale, AI, and beyond Moore's Law capabilities such as QIS.

Activities and Explanation of Changes

(dollars in thousands)

FY 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
Mathematical, Computational, and Computer Sciences Research	**$259,865**	**$292,962**	**+$33,097**
Applied Mathematics Research	$48,570	$51,048	+$2,478
Funding expands 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 Request will continue to 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. Request also supports planning for the transition of critical Applied Math efforts from the Exascale Computing Project (ECP) into core research areas.	Funding will continue support of core research efforts and incorporating ECP efforts into core research areas.
Computer Sciences Research	$46,827	$49,773	+$2,946
Funding continues 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 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 continue long-term basic research efforts that explore and prepare for emerging technologies, such as quantum networking, specialized and heterogeneous hardware and accelerators, and quantum information systems. Request also supports planning for the sustainability of software from the ECP.	Increased funding will support core investments and continue long-term basic research initiative efforts.

(dollars in thousands)

FY 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
Computational Partnerships	$76,194	$86,029	+$9,835
Funding continues 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 Request will continue support for the SciDAC Institutes and partnerships with SC and DOE applications. Partnerships on scientific data, AI, microelectronics research, Quantum Information Science, and an integrated computational and data infrastructure for science will continue. The partnership with NIH will increase to 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. New efforts will focus on enabling widespread use of DOE high performance computing resources by Federal agencies in support of emergency preparedness and response. The Request also supports planning for the transition of mission critical Exascale Computing Project applications into SciDAC.	The increase will support interagency efforts to maximize the public benefit of DOE infrastructure and accelerate development of AI techniques for science. The increase will also support investments in capabilities that allow DOE to rapidly respond to research needs in support of national emergencies. This includes partnering with key agencies to understand their simulation and modeling capabilities, data management and curation needs, and identify and bridge gaps necessary for DOE to provide resources on short notice.

(dollars in thousands)

FY 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
Advanced Computing Research	$88,274	$106,112	+$17,838
Funding continues to support quantum testbed efforts, with emphasis on partnerships with the new QIS centers. 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. Funding under this activity continues to support small investments in REP for cybersecurity and testbeds for advanced microelectronics research. In addition, funding provides support for the CSGF fellowship at $10,000,000, in partnership with NNSA. The goal of CSGF is to increase availability of a trained workforce for exascale, AI, and beyond Moore's Law capabilities such as QIS.		The Request will continue to support the National QIS Research Centers and quantum testbed efforts. Building on basic research in quantum information networks, this activity will continue to support early-stage research associated with the first steps to establishing a dedicated Quantum network. The Request allows REP to explore new strategic investments in emerging technologies and microelectronics. Small investments in cyber security will continue. The Request will increase support for the CSGF fellowship, in partnership with NNSA, to support increased tuition costs, to increase the number of fellows focused on emerging technologies, and to expand the participation of groups, fields, and institutions that are under-represented in high end computational science. The goal of CSGF is to increase availability of a trained workforce for exascale, AI, and beyond Moore's Law capabilities such as QIS. The Request also supports the RENEW initiative to provide undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem to expand the pipeline for ASCR research and facilities workforce needs.	The increase will support early stage research for a quantum network, support for the National QIS Research Centers, selected in FY 2020, and user access to quantum resources at the quantum testbeds and other quantum computers through the quantum computing user program at ORNL. Funding will also support an increase in support for the CSGF and the RENEW initiative.

Note: Funding for the subprogram above, includes 3.65% of research and development (R&D) funding for the Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Programs.

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.

ASCR regularly gathers requirements from the other SC research programs through formal processes, including workshops and technical reviews, to inform upgrade plans and user programs. 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 ASCR HPC facilities computing resources to users follows the merit review public-access model used by other SC scientific user facilities. The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program provides access to the LCFs and the ASCR Leadership Computing Challenge (ALCC) program provides an allocation path for critical DOE mission applications and to broaden the community of users on the LCFs and NERSC.

ASCR HPC facilities received CARES Act funding in FY 2020 to support COVID-19 research teams through the HPC Consortium. DOE deployed customized hardware to each HPC facility to better address specific COVID-19 research needs and dozens of projects used millions of hours of compute time to provide insights including how the virus infects cells, exploration of treatment options, understanding variation in patient outcomes, high throughput drug candidate screening, epidemiology and public health surveillance, and advanced data analytics. Several of these research teams were finalists for the special COVID-19 Gordon Bell Prize. Once the pandemic is over, this hardware will be integrated into the HPC resources available for peer-reviewed, competitive research and will continue to advance biological and medical research.

<u>High Performance Production Computing</u>

This activity supports NERSC at LBNL to deliver high-end production computing services for the SC research community. More than 8,000 computational scientists conducting about 700 projects use NERSC annually to perform scientific research across a wide range of disciplines including astrophysics, chemistry, earth systems modeling, materials science, engineering, 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, approximately one-third in DOE laboratories, and other users from government laboratories, non-profits, small businesses, 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 NERSC's HPC resources.

NERSC currently operates the 30 petaflops (pf) Intel/Cray NERSC-8 system (Cori), as well as the 125 pf HPE/AMD/NVIDIA NERSC-9 system (Perlmutter), that comes online in FY 2021. 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. In FY 2022, NERSC will initiate preparations for a NERSC-10 upgrade that is needed to meet SC's growing computational needs.

Current ASCR HPC resources and facilities are primarily designed to efficiently execute large-scale simulation data analysis applications on premises, with a focus on minimizing users' wait-times in batch queues. With the rate and volume of data generated at SC scientific user facilities growing exponentially, and the significant interest in coupling these facilities to HPC in real-time to drive scientific discovery, ASCR is exploring the resource requirements to enable HPC for a broadening set of mission essential computational science workflows. In FY 2022, ASCR will design a state-of-the-art Scientific High Performance Data Facility focused on the unique challenges and opportunities for data-intensive applications workflows and near real-time computing needed to support the explosion of SC scientific data that will serve as the anchor for the Integrated Computational and Data Infrastructure Initiative. To provide geographic diversity and operational resiliency, this facility will be located on the East Coast.

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 strength of the ASCR facilities 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 OLCF-4 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. The OLCF will undertake final acceptance and operations of an HPE-Cray/AMD exascale system (Frontier), which was deployed in calendar year 2021.

The Argonne Leadership Computing Facility (ALCF) at ANL operates the 8.5 pf Intel/Cray ALCF-2 system (Theta) and testbeds such as Polaris (A-19) to prepare their users and SC-ECP applications and software technology for the ALCF-3 upgrade, to be known as Aurora. Aurora will be an exascale system when deployed in calendar year 2022 and is being designed by Intel/HPE-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 new sustainable materials to capture and convert solar energy; developing new solid-state energy storage systems; 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 manage enterprise risk and foster diverse capabilities that provide the Nation's HPC user community with the most effective resources.

The demand for 2021 INCITE allocations at the LCFs outpaced the available resources by more than a factor of three. Demand for 2020-2021 ALCC allocations outpaced resources by more than a factor of five.

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 executing a complete upgrade of its backbone network (the ESnet-6 upgrade).

Activities and Explanation of Changes

(dollars in thousands)

FY 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
High Performance Computing and Network Facilities	**$586,190**	**$618,038**	**+$31,848**
High Performance Production Computing	$113,786	$115,963	+$2,177
Funding supports 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.		The Request will support operations at the NERSC facility, including user support, power, space, system leases, and staff. The Request will also support full operations of Perlmutter, which will be installed and accepted in FY 2021.	Funding will support the operations of the Cori and Perlmutter systems as well as development of early-stage designs for the NERSC-10 upgrade. In addition, funding will also support the design of a Scientific High-Performance Data Facility located on the East Coast.
Leadership Computing Facilities	$381,075	$408,113	+$27,038
Funding supports 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.		The Request will support operations at the LCF facilities at ANL and ORNL, including user support, power, space, system leases, early access systems and testbeds, and operations staff. The Request also will support deployment, acceptance testing, early science and operations of exascale systems at OLCF and ALCF.	Funding will support the calendar year 2021–2022 deployment schedule for both LCF upgrades to the exascale.
Leadership Computing Facility at ANL	*$152,955*	*$159,047*	*+$6,092*
Funding continues 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.		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 calendar year 2022.	Funding will continue operations at the ALCF and support final preparations, deployment and early science access of the ALCF-3 exascale system.

(dollars in thousands)

FY 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
Leadership Computing Facility at ORNL	*$228,120*	*$249,066*	*+$20,946*
Funding continues 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 Request will continue support for the operation and competitive allocation of the Summit system and other testbeds. The OLCF will complete acceptance of the OLCF-5 exascale system, Frontier, which was deployed in calendar year 2021 and will provide access for early science applications and the Exascale Computing Project.	The increase in funding will support the operations of the exascale system deployed in calendar year 2021.
High Performance Network Facilities and Testbeds	$91,329	$93,962	+$2,633
Funding supports 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 on dark fiber with new equipment, increased capacity, and an advanced 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 ESnet-6 upgrade in accordance with the project baseline as well as continued operation of ESnet-5.

Note: Funding for the subprogram above, includes 3.65% of research and development (R&D) funding for the Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Programs.

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 (LCFs) 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 its major computing projects, including the LCFs at ANL and ORNL, and NERSC at LBNL.

SC conducts overall project management for the SC-ECP via a Project Office established at ORNL because of its considerable expertise in developing computational science and engineering applications and in managing HPC facilities, both for the Department and for other federal agencies; and its 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 2022 Request includes $129,000,000 for the SC-ECP. These funds will support test and deployment investments in the ECP technical focus areas—application development, software technology and hardware and integration—via testing and development on early access exascale hardware and exascale systems. 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 2021 Enacted		FY 2022 Request	Explanation of Changes FY 2022 Request vs FY 2021 Enacted
Construction	**$168,945**	**$129,000**	**-$39,945**
17-SC-20, SC Exascale Computing Project	$168,945	$129,000	-$39,945
Funding supports 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.		The Request will support project management; co-design activities between application and the software stack; and integration between SC-ECP and the LCFs to provide continuous integration and testing of the ECP funded applications and software on exascale testbed.	The funding decrease reflects the completion of research and development activities as the project moves into implementation of software technology and execution of application challenge problems on the actual exascale systems.

Advanced Scientific Computing Research
Capital Summary

(dollars in thousands)

	Total	Prior Years	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Capital Operating Expenses						
Capital Equipment	N/A	N/A	5,000	31,809	5,000	-26,809
Total, Capital Operating Expenses	**N/A**	**N/A**	**5,000**	**31,809**	**5,000**	**-26,809**

Capital Equipment

(dollars in thousands)

	Total	Prior Years	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Capital Equipment						
Total, Non-MIE Capital Equipment	N/A	N/A	5,000	31,809	5,000	-26,809
Total, Capital Equipment	**N/A**	**N/A**	**5,000**	**31,809**	**5,000**	**-26,809**

Advanced Scientific Computing Research
Funding Summary

(dollars in thousands)

	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Research	405,000	428,810	421,962	-6,848
Facility Operations	575,000	586,190	618,038	+31,848
Total, Advanced Scientific Computing Research	**980,000**	**1,015,000**	**1,040,000**	**+25,000**

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.

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 2020 Enacted	FY 2020 Current	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Scientific User Facilities - Type A					
National Energy Research Scientific Computing Center	**110,000**	**109,224**	**111,245**	**115,963**	**+4,718**
Number of Users	7,500	8,329	8,329	8,500	+171
Achieved Operating Hours	–	8,293	–	–	–
Planned Operating Hours	8,585	8,585	8,585	8,585	–
Optimal Hours	8,585	8,585	8,585	8,585	–
Percent of Optimal Hours	100.0%	100.0%	100.0%	100.0%	–
Unscheduled Down Time Hours	–	292	–	–	–
Argonne Leadership Computing Facility	**150,000**	**150,000**	**153,818**	**159,047**	**+5,229**
Number of Users	950	1,174	1,174	1,300	+126
Achieved Operating Hours	–	6,923	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Optimal Hours	7,008	7,008	7,008	7,008	–
Percent of Optimal Hours	100.0%	100.0%	100.0%	100.0%	–
Unscheduled Down Time Hours	–	84	–	–	–
Oak Ridge Leadership Computing Facility	**225,000**	**225,000**	**229,915**	**249,066**	**+19,151**
Number of Users	1,450	1,546	1,546	1,500	-46
Achieved Operating Hours	–	6,980	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Optimal Hours	7,008	7,008	7,008	7,008	–
Percent of Optimal Hours	100.0%	100.0%	100.0%	100.0%	–
Unscheduled Down Time Hours	–	28	–	–	–

(dollars in thousands)

	FY 2020 Enacted	FY 2020 Current	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Energy Sciences Network	**90,000**	**90,000**	**91,212**	**93,962**	**+2,750**
Achieved Operating Hours	–	8,760	–	–	–
Planned Operating Hours	8,760	8,760	8,760	8,760	–
Optimal Hours	8,760	8,760	8,760	8,760	–
Percent of Optimal Hours	100.0%	100.0%	100.0%	100.0%	–
Total, Facilities	**575,000**	**574,224**	**586,190**	**618,038**	**+31,848**
Number of Users	9,900	11,049	11,049	11,300	+251
Achieved Operating Hours	–	30,956	–	–	–
Planned Operating Hours	31,361	31,361	31,361	31,361	–
Optimal Hours	31,361	31,361	31,361	31,361	–
Unscheduled Down Time Hours	–	404	–	–	–

Note: Achieved Operating Hours and Unscheduled Downtime Hours will only be reflected in the Congressional budget cycle which provides actuals.

Advanced Scientific Computing Research
Scientific Employment

	FY 2020 Enacted	FY 2021 Enacted	FY 2022 Request	FY 2022 Request vs FY 2021 Enacted
Number of Permanent Ph.Ds (FTEs)	683	814	825	+11
Number of Postdoctoral Associates (FTEs)	331	349	365	+16
Number of Graduate Students (FTEs)	438	520	535	+15
Number of Other Scientific Employment (FTEs)	212	217	220	+3

Note: Other Scientific Employment (FTEs) includes technicians, engineers, computer professionals and other support staff.

17-SC-20, SC Exascale Computing Project

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

Summary

The FY 2022 Request for the SC Exascale Computing Project (SC-ECP) is $129,000,000. The Total Estimated Cost (TEC) range for this project is between $1,500,000,000 and $1,900,000,000. The Total Project Cost (TPC) of the SC portion of ECP is $1,326,200,000 with the total combined SC and NNSA TPC of $1,812,300,000.

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.

Other activities included in the ECI but not the Office of Science-Exascale Computing Project (SC-ECP) include $275,000,000 in FY 2022 to support the final site preparations and the installation of the exascale system at the Argonne Leadership Computing Facility (ALCF) and the final acceptance of the exascale system deployed at the Oak Ridge Leadership Computing Facilities (OLCF) in calendar year 2021. Supporting parallel development at both Leadership Computing Facilities (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.

Significant Changes

This project was initiated in FY 2017. The FY 2022 Request supports investments in the ECP technical focus areas—application development, software technology and hardware and integration—to initiate test and deployment a capable exascale software ecosystem on exascale-capable systems deployed in the calendar year 2021 and 2022. Funding decreases in FY 2022 to reflect the completion of research and development activities as the project focus moves more to execution and implementation.

The most recent DOE Order 413.3B approved Critical Decision (CD) is CD-2/3 Approve Performance Baseline, where definitive scope, schedule and cost baselines have been developed and the project is ready for implementation. The project achieved CD-2/3 on February 25, 2020. The Total Project Cost (TPC) of the SC portion of ECP is $1,326,200,000 with the total combined SC and NNSA TPC of $1,812,300,000.

Critical Milestone History

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	CD-2	Final Design Complete	D&D Complete	CD-4
FY 2017	3Q FY 2016	TBD	TBD	TBD	TBD	N/A	TBD
FY 2018	7/28/16	2Q FY 2019	1/3/17	4Q FY 2019	3Q FY 2019	N/A	4Q FY 2023
FY 2019	7/28/16	2Q FY 2019	1/3/17	4Q FY 2019	3Q FY 2019	N/A	4Q FY 2023
FY 2020	7/28/16	2Q FY 2019	1/3/17	1Q FY 2020	3Q FY 2019	N/A	4Q FY 2023
FY 2021	7/28/16	3/22/16	1/3/17	2Q FY 2020	6/6/19	N/A	4Q FY 2024
FY 2022	7/28/16	3/22/16	1/3/17	2/25/20	6/6/19	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-3 – 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	CD-3B
FY 2017	TBD	TBD	TBD
FY 2018	4Q FY 2019	1/3/17	4Q FY 2019
FY 2019	4Q FY 2019	1/3/17	4Q FY 2019
FY 2020	1Q FY 2020	1/3/17	1Q FY 2020
FY 2021	1Q FY 2020	1/3/17	2Q FY 2020
FY 2022	2/25/20	1/3/17	2/25/20

CD-3A – Approve Long Lead Time Procurements
CD-3B – Approve Remaining Construction Activities

Project Cost History

(dollars in thousands)

Fiscal Year	TEC, Design	TEC, Construction	TEC, Total	OPC, Except D&D	OPC, D&D	OPC, Total	TPC
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
FY 2022	N/A	700,843	700,843	625,363	N/A	625,363	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

In 2015, the National Strategic Computing Initiative was established to maximize the benefits of HPC for U.S. economic competitiveness, scientific discovery, and national security. Within that initiative DOE, represented by a partnership between SC and NNSA, has the responsibility for executing a joint program focused on advanced simulation through an exascale–capable computing program, which will emphasize sustained performance and analytic computing to advance DOE missions. The objectives and the associated scientific challenges define a mission need for a computing capability of 2 – 10 ExaFLOPS (2 billion billion floating-point operations per second) in the early to mid-2020s. In FY 2017, SC initiated the 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.

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 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 their KPP-1 stretch goal.
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 stretch goal
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

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Construction (TEC)			
FY 2019	187,696	187,696	–
FY 2020	174,735	174,735	105,440
FY 2021	160,412	160,412	336,480
FY 2022	115,000	115,000	154,909
Outyears	63,000	63,000	104,014
Total, Construction (TEC)	**700,843**	**700,843**	**700,843**
Total Estimated Cost (TEC)			
FY 2019	187,696	187,696	–
FY 2020	174,735	174,735	105,440
FY 2021	160,412	160,412	336,480
FY 2022	115,000	115,000	154,909
Outyears	63,000	63,000	104,014
Total, TEC	**700,843**	**700,843**	**700,843**

[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.

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Other Project Cost (OPC)			
FY 2016	146,820	146,820	5,741
FY 2017	164,000	164,000	83,698
FY 2018	205,000	205,000	170,898
FY 2019	45,010	45,010	220,565
FY 2020	14,000	14,000	93,203
FY 2021	8,533	8,533	9,258
FY 2022	14,000	14,000	14,000
Outyears	28,000	28,000	28,000
Total, OPC	**625,363**	**625,363**	**625,363**

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Project Cost (TPC)			
FY 2016	146,820	146,820	5,741
FY 2017	164,000	164,000	83,698
FY 2018	205,000	205,000	170,898
FY 2019	232,706	232,706	220,565
FY 2020	188,735	188,735	198,643
FY 2021	168,945	168,945	345,738
FY 2022	129,000	129,000	168,909
Outyears	91,000	91,000	132,014
Total, TPC	**1,326,206**	**1,326,206**	**1,326,206**

4. Details of Project Cost Estimate

The SC-ECP was baselined at CD-2. The Total Project Cost for the SC-ECP is represented in the table below.

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Application Development (TEC)	347,289	236,759	346,360
Production Ready Software	228,472	142,661	217,290
Hardware Partnership	125,082	128,260	131,726
Total, Other (TEC)	**700,843**	**507,680**	**695,376**
Total, TEC	**700,843**	**507,680**	**695,376**
Contingency, TEC	***N/A***	***N/A***	***N/A***
Other Project Cost (OPC)			
Planning Project Management	89,688	102,595	89,688
Application Development (OPC)	221,050	333,568	221,050
Software Research	118,517	199,754	118,517
Hardware Research	196,108	182,609	201,575
Total, Except D&D (OPC)	**625,363**	**818,526**	**630,830**
Total, OPC	**625,363**	**818,526**	**630,830**
Contingency, OPC	***N/A***	*N/A*	***N/A***
Total, TPC	**1,326,206**	**1,326,206**	**1,326,206**
Total, Contingency (TEC+OPC)	***N/A***	***N/A***	***N/A***

5. Schedule of Appropriations Requests

(dollars in thousands)

Request Year	Type	Prior Years[a]	FY 2020	FY 2021	FY 2022	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	707,524	14,000	—	—	42,000	763,524
	TPC	707,524	189,000	—	—	257,000	1,153,524

[a] 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.

(dollars in thousands)

Request Year	Type	Prior Years[a]	FY 2020	FY 2021	FY 2022	Outyears	Total
FY 2019	TEC	—	174,735	—	—	252,000	426,735
	OPC	751,230	14,000	—	—	42,000	807,230
	TPC	751,230	188,735	—	—	294,000	1,233,965
FY 2020	TEC	—	174,735	—	—	252,000	426,735
	OPC	759,650	14,000	—	—	56,000	829,650
	TPC	759,650	188,735	—	—	308,000	1,256,385
FY 2021	TEC	—	174,735	154,945	—	178,000	507,680
	OPC	748,526	14,000	14,000	—	42,000	818,526
	TPC	748,526	188,735	168,945	—	220,000	1,326,206
FY 2022	TEC	187,696	174,735	160,412	115,000	63,000	700,843
	OPC	560,830	14,000	8,533	14,000	28,000	625,363
	TPC	748,526	188,735	168,945	129,000	91,000	1,326,206

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 2022 Budget Request, $275,000,000, is included in the ALCF and OLCF 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	7 years
Expected Future Start of D&D of this capital asset	2029

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.

[a] 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.