

Advanced Scientific Computing Research

Overview

The mission of the Advanced Scientific Computing Research (ASCR) program 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 applications and Department of Energy (DOE) mission needs. ASCR's partnerships, including new efforts with the applied technology offices, activities to broaden participation of under-served communities, 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 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, and stockpile stewardship without testing. Computational science allows researchers 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 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 strengthen our clean energy economy and ensure all Americans share the benefits from those investments. 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 scales 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 and technology development and deployment.

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 redefine and construct the foundation for a new 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, the applied technology offices, other agencies, 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 at the exascale and beyond to accelerate progress in delivering a clean energy future, understanding and addressing climate change, broadening the impact of our investments in science, and increasing the competitive advantage of U.S. industry.

Highlights of the FY 2023 Request

The FY 2023 Request of \$1,068.7 million for ASCR will strengthen U.S. leadership in strategic computing with operation and allocation of the Nation's first exascale computing system and testing of a second system, broadening the foundations of AI and QIS, and expanding the infrastructure and partnerships that enables data-driven science and technology—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 Computing Project (ECP) deployment, the Request prioritizes foundational research in Applied Mathematics and Computer Science and the transition of critical technologies from the ECP. Investments will continue to emphasize foundational research to address the combined challenges of increasingly heterogeneous architectures and the changing ways in which HPC systems are used, the development of new scalable energy efficient algorithms and software, directed basic research to address specific challenges for the new Energy Earthshot Research Centers, and incorporating AI and ML into simulations and data intensive applications while increasing greater connectivity with distributed resources, including other SC user facilities. The Computational Partnerships activity will continue to infuse the latest developments in applied math and computer science into strategic applications, including areas such as 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. The Request increases support for ASCR's Computational Partnerships to support SC initiatives; extends SciDAC partnerships to DOE's applied technology offices and mission critical ECP applications, including interagency partnerships to actively build workflows that ensure rapid and robust response to emerging pandemic, biothreat, and public health emergencies; and expands the SciDAC Institutes to fully incorporate exascale software and libraries. The Request also sustains increased 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 Research Centers (NQISRCs), quantum internet, and testbeds. ASCR will continue to partner with the other SC programs to support the multi-disciplinary NQISRCs. 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 regional quantum testbeds, which provide researchers with access to novel, early-stage quantum computing and networking resources and services, and basic research in quantum information will continue. In FY 2023, ASCR will begin to enable the sustainability of critical ECP software for use on emerging technology testbeds.
- ASCR increases support for 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. Science and Technology (S&T) ecosystem.
- The Funding for Accelerated, Inclusive Research (FAIR) initiative will provide focused investment on enhancing research on clean energy, climate, and related topics at minority serving institutions, including attention to underserved and environmental justice regions. The activities will improve the capability of MSIs to perform and propose competitive research and will build beneficial relationships between MSIs and DOE national laboratories and facilities.
- ASCR will also participate in the Accelerate initiative through support for scientific research that will accelerate the transition of science advances to energy technologies. These new efforts will drive scientific discovery toward sustainable production of new technologies across the innovation continuum, to provide experiences in working across this continuum for the workforce needed for industries of the future, and to meet the nation's needs for abundant clean energy, a sustainable environment, and national security.

- Energy Earthshot Research Centers (EERCs), a new modality of research to be launched in FY 2023, building on the success of SC's Energy Frontier Research Centers (EFRCs) and the SciDAC program, will bring together multi-investigator, multi-disciplinary teams to perform energy-relevant research with a scope and complexity beyond what is possible in standard single-investigator or small-group awards. Beyond complementing and expanding the scope of the EFRCs and SciDAC, the EERCs will address the research challenges at the interface between currently supported basic research, applied research, and development activities, with support from both SC and the applied technology offices.

Facility Operations

- FY 2023 marks the beginning of the exascale era for the U.S. research community with full operations and competitive allocation of the Nation's first exascale computing system at the Oak Ridge Leadership Computing Facility (OLCF), a system called Frontier that was deployed in calendar year 2021, and acceptance and operations for early science and the ECP applications on a second exascale computing system at the Argonne Leadership Computing Facility (ALCF), a system called Aurora. The Request 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 operations costs at the Leadership Computing Facilities, at 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 completion of the Department's Exascale Computing Initiative (ECI), which includes the SC-Exascale Computing Project (SC-ECP). In addition, the Request supports the NERSC-10 upgrade, planned to start in FY 2022, including site preparations and long lead procurements, to address rising demand for production computing across the SC programs.
- 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. In addition, the diversity of data- and compute-intensive research workflows is expanding rapidly. In FY 2023, ASCR will continue planning for a new High Performance Data Facility (HPDF), to satisfy the unique requirements of state-of-the-art real-time experimental/observational workflows. The NERSC-10 and HPDF projects will each drive unique technological innovation in system architectures and services beyond what is available in the commercial cloud and will inform planning for future upgrades at the LCFs.

Projects

- The ASCR FY 2023 Request includes \$227.0 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 delivered in calendar year 2021 and a second in 2022 to address national needs. Of this effort, \$150 million will go to the ALCF to deploy and operate Aurora and testbeds in support of the ECP project teams.

**Advanced Scientific Computing Research
FY 2023 Research Initiatives**

Advanced Scientific Computing Research supports the following FY 2023 Research Initiatives.

	(dollars in thousands)			
	FY 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Accelerate Innovations in Emerging Technologies	–	–	5,000	+5,000
Advanced Computing	–	–	15,332	+15,332
Artificial Intelligence and Machine Learning	56,866	58,820	73,000	+16,134
Biopreparedness Research Virtual Environment (BRaVE)	–	–	11,183	+11,183
Exascale Computing	438,945	404,000	227,000	-211,945
Funding for Accelerated, Inclusive Research (FAIR)	–	–	4,073	+4,073
Integrated Computational & Data Infrastructure	11,974	14,321	–	-11,974
Microelectronics	5,182	5,183	5,183	+1
Quantum Information Science	98,402	100,695	101,194	+2,792
Reaching a New Energy Sciences Workforce (RENEW)	–	–	10,000	+10,000
SC Energy Earthshot	–	–	50,000	+50,000
Total, Research Initiatives	611,369	583,019	501,965	-109,404

Note:

- The Integrated Computational and Data Initiative is rolled into Advanced Computing Initiative in FY 2023.

**Advanced Scientific Computing Research
Funding**

(dollars in thousands)

	FY 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Advanced Scientific Computing Research				
Applied Mathematics Research	48,570	48,570	71,938	+23,368
Computer Sciences Research	46,827	49,388	70,326	+23,499
Computational Partnerships	76,194	73,892	97,861	+21,667
Advanced Computing Research	88,274	96,112	113,598	+25,324
Energy Earthshot Research Centers	—	—	25,000	+25,000
Total, Mathematical, Computational, and Computer Sciences Research	259,865	267,962	378,723	+118,858
High Performance Production Computing	113,786	115,963	115,033	+1,247
Leadership Computing Facilities	381,075	408,113	407,772	+26,697
High Performance Network Facilities and Testbeds	91,329	93,962	90,213	-1,116
Total, High Performance Computing and Network Facilities	586,190	618,038	613,018	+26,828
17-SC-20 SC Exascale Computing Project	168,945	129,000	77,000	-91,945
Subtotal, Advanced Scientific Computing Research	1,015,000	1,015,000	1,068,741	+53,741
Total, Advanced Scientific Computing Research	1,015,000	1,015,000	1,068,741	+53,741

SBIR/STTR funding:

- FY 2021 Enacted: SBIR \$25,736,000 and STTR \$3,620,000
- FY 2022 Annualized CR: SBIR \$27,495,000 and STTR \$3,869,000
- FY 2023 Request: SBIR \$30,775,000 and STTR \$4,327,000

**Advanced Scientific Computing Research
Explanation of Major Changes**

(dollars in thousands)

FY 2023 Request vs FY 2021 Enacted

+\$118,858

Mathematical, Computational, and Computer Sciences Research

The Computer Science and Applied Mathematics activities will: continue to increase their efforts on foundational research and long-term basic research efforts that explore and prepare for emerging technologies such as quantum networking and computing; begin to transition critical technologies from the Exascale Computing Project into core research efforts; develop new scalable energy efficient algorithms and software; address specific challenges for the Energy Earthshot Research Centers; and address the challenges of data intensive science and the development of critical tools, including AI/ML, to enable an integrated computational and data infrastructure. Computational Partnerships will increase to support SC initiatives; extend SciDAC partnerships to DOE's Applied Energy programs and mission critical ECP applications, including interagency partnerships and emergency preparedness, and to expand the SciDAC Institutes to fully incorporate exascale software and libraries. The Advanced Computing Research activity continues to robustly support the National QIS Research Centers, quantum testbeds, and regional quantum networking testbeds, in close coordination with the other SC programs, to expand user access to quantum resources. Increased funding will enable the sustainability of critical ECP software for use on emerging technology testbeds. This subprogram also increases support for 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 and sustains increased support for the Computational Sciences Graduate Fellowship. The subprogram also supports ASCR's participation in SC's new Accelerate and FAIR initiatives to expand participation, accelerate innovation, reduce impacts from climate change, and advance clean energy technologies and infrastructure.

High Performance Computing and Network Facilities

The OLCF will provide full operations and competitive allocation of the Nation's first exascale computing system, Frontier, deployed in calendar year 2021. The ALCF will complete acceptance testing and early science/ECP access to the Aurora exascale computing system, deployed in calendar year 2022. Both facilities will provide testbed resources to explore new technologies. In addition, funding supports operation of the 125 petaflop NERSC-9 Perlmutter system, site preparations and long-lead procurements for NERSC-10, and the ESnet-6 upgrade in accordance with the project baselines. To address the significant growth in the rate and volume of data from SC scientific user facilities, ASCR also initiated planning for a new High Performance Data Facility (HPDF) initiated in the FY 2022 Request to satisfy the unique requirements of state-of-the-art real-time experimental/observational workflows to support the explosion of data and also serve as the anchor for the integrated computational and data infrastructure efforts. Funding for all facilities supports operations, including power, equipment, staffing, testbeds, lease payments, and planning for future upgrades.

+\$26,828

(dollars in thousands)

FY 2023 Request vs FY 2021 Enacted

-\$91,945

Exascale Computing

The FY 2023 Request will support efforts to deploy SC-ECP applications and ecosystem on both exascale architectures in partnership with the ASCR facilities. The decrease represents a shift in focus within the project as it matures beyond delivery to the execution of applications' challenge problems and implementation of software technologies to meet the Key Performance Parameters (KPPs) on the exascale systems delivered in 2021 and 2022.

Total, Advanced Scientific Computing Research

+\$53,741

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, Scientific Data, Large Scale Networking, AI, High End 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 National Nuclear Security Administration (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 Networking and Information Technology R&D Subcommittee of the NSTC Committee on Technology, ASCR also coordinates with programs across the Federal Government. In FY 2023, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Supercomputing Versus COVID-19 – Round Two, Understanding COVID Variants

The SARS-CoV-2 pandemic has entered a new phase with the emergence of variants of concern (VOC) that are more contagious and could undermine the protection of vaccines. Substitutions in the spike protein have been identified with new variants but do not fully explain the success of fast-spreading variants, like the Delta and Omicron variants. A research team at ORNL used a computational systems biology approach to process more than 900,000 SARS-CoV-2 genomes and map spatiotemporal relationships, revealing other critical attributes of successful variants. Comparisons to earlier dominant mutations and protein structural analyses indicate that the increased transmission is promoted by the combination of functionally complementary mutations in both the spikes and in other regions of the SARS-CoV-2 proteome. They found that the currently known VOCs have common mutations in proteins involved in immune-antagonism and replication performance, suggesting a convergent evolution of the virus. Critically, they found that VOCs often occur with a sudden doubling of the number of mutations in that strain, which indicates recombination events—where two different strains infect the same person and, in the process of replicating, merge their genetic material—that facilitate the combination of mutations in spike and other proteins that together lead to new, more infectious and increasingly immune escaping variants. This indicates that extensive community distribution of numerous SARS-CoV-2 variants increases the probability of future recombination events, further accelerating the evolution of the virus.

Toward an Integrated Scientific Data Infrastructure - Addressing Urgent Challenges in Real Time

Scientists working remotely from Turkey leveraged the National Energy Research Scientific Computing Center (NERSC) and the Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory, via the ESnet, to capture detailed images of the structure of the SARS-CoV-2 virus, focusing on two of the viral proteases—enzymes that make the virus's life cycle possible—and how to keep them from functioning. By understanding the molecular structure of the proteases, researchers can identify proteins that bind to them and interfere with their role in viral reproduction. The goal is to inhibit these two enzymes with chemical compounds contributed from the COVID-19 Moonshot initiative that may eventually lead to antiviral treatments in humans, a key step toward next generation treatments for COVID-19. During and after the experiments, the team utilized NERSC's Cori supercomputer and ESnet's high-speed optical network to process data and provide results in real time, allowing the researchers to monitor the experiment, begin analysis, and make changes as necessary. This enables the study of small differences in atomic structure in near real time, even while working remotely, to guide decision-making during the experiment. Analysis of the findings from these experiments is ongoing but the work is also shaping the development of the ASCR scientific data facility concept, exposing challenges and identifying areas where more work is needed.

ASCR Inside: Taking AI to Extreme Scales for Science

Scientists at Los Alamos National Laboratory and the SLAC National Accelerator Laboratory, working in collaboration with NVIDIA, Facebook, and other industrial and academic teams, leveraged ASCR's long-standing investments in the Legion data-centric parallel-programming system to create the FlexFlow large-scale machine-learning toolkit. This toolkit allows practitioners with workflows programmed with widely-used machine-learning frameworks and interfaces, such as Keras/TensorFlow, PyTorch, and ONNX, to utilize large-scale resources to train machine-learning models of unprecedented size and faster than ever. Efficiently using large-scale resources enabled a 15x reduction in model-training times, bringing a key benchmark workflow from 18 hours to 1.2 hours. The increased turn-around time was enabled by scaling the workflow

to use over 750 Graphics Processing Units (GPUs), a feat made possible using the FlexFlow framework build on Legion, an R&D 100 Award winner in 2020. The workflow acceleration contributed to urgent scientific work aimed at discovering new drugs to treat cancer and COVID-19.

Getting Real About Uncertainty in Sea Level Rise

A team of DOE-supported ice sheet and climate modeling scientists contributed to improved quantitative estimates for the range of future sea level rise expected from melting glaciers and ice sheets, as recently reported on in the journal *Nature*.^a This study was also cited in Chapter 9 of the 6th assessment report from the International Panel on Climate Change (IPCC), "Ocean, Cryosphere and Sea Level Change". For this study, the team of 84 international researchers used NERSC to run the largest and most sophisticated set of climate and land ice models to date, combining nearly 900 simulations from 38 international modeling groups to improve not only the median projections of future sea level rise but also estimates for the associated uncertainties. These multi-model ensembles were combined via statistical emulation to build probabilistic projections of future sea level rise from all sources of land ice. The projections show that limiting global warming to 1.5°C above pre-industrial temperatures would cut projected 21st century sea level rise from land ice in half, relative to currently pledged emissions reductions, from approximately 25 cm to 13 cm in the best-case scenario. However, under the worst-case scenario, with much more melting than snowfall in the Antarctic, ice losses there could be five times larger, increasing the median land ice contribution to 42 cm of sea level rise under current policies and pledges, with a 5 percent chance of sea level rise exceeding 50 cm even under 1.5°C warming. Results from this study confirm that Antarctica remains a critical focus for reducing future uncertainty in sea level rise; due to substantial uncertainty in how strongly warm ocean waters will erode floating parts of the ice sheet from beneath, a process that now is the focus of DOE's Energy Exascale Earth System Model.

Supercomputing Powers Energy Savings

Residential and commercial buildings consume nearly three-quarters of U.S. electricity—during peak hours, the share reaches 80 percent. The annual energy bill is nearly \$412 billion. Simulating that energy use on a broad scale can help identify ways to reduce it, cutting greenhouse gas emissions in the process. Researchers at Oak Ridge National Laboratory used the Argonne Leadership Computing Facility for data-intensive simulations of a "digital twin" of the more than 178,000 buildings in Chattanooga, Tennessee. They studied how different energy conservation measures might result in cost savings. The AutoBEM simulation of Chattanooga buildings, which also uses EnergyPlus and OpenStudio, two DOE tools for modeling buildings energy use, found that 99 percent of buildings would realize energy savings from employing the existing energy efficiency technologies evaluated. The effort is part of a larger goal to model all of the Nation's 125 million buildings. A conservation measure of the impact of simple changes like improved HVAC efficiency, space sealing, insulation or lighting could have the potential to offset 500 to 3,000 pounds of carbon dioxide per building, the researchers concluded.

Delivering the Exascale Ecosystem

The Exascale Computing Project (ECP) is building a comprehensive software ecosystem consisting of more than 20 applications and 80 software packages that use more than 10 compilers and 10 programming models on a range of target hardware architectures (from laptops to exascale). The legacy of these efforts will be a computational ecosystem that accelerates U.S. capabilities in scientific simulation and artificial intelligence (AI), unlocking the potential of Exascale computers and preparing us for future systems that will build on our legacy. The complexity of this ecosystem, with over one million combinations, is being managed and simplified for the entire scientific and AI user community through the creation of ECP's Extreme-Scale Scientific Software Stack (E4S). E4S (<https://e4s.io>) is lowering the barrier to entry for users and developers in the DOE and other U.S. government agencies, industries, and universities. The E4S has aggressively evolved, providing 80 distinct turn-key HPC and AI products grouped into thematic software development kits (SDKs). Application codes that build on top of E4S benefit from guaranteed version compatibility, access to the latest stable features integrated into each quarterly release, and advanced build environment features that can improve build times by a factor of ten. The last E4S release contains support for GPU architectures and is installed on DOE and NSF pre-exascale systems, enabling the portability promise that is central to the success of future U.S. supercomputing.

^a Edwards, T. L. and 82 others. *Nature*, 593(7857), 74–82, <https://www.nature.com/articles/s41586-021-03302-y>

Physics-Informed Machine Learning

The accelerating pace of observational data from experiments far outpaces our ability to understand it. Despite the promise and some preliminary success, most machine learning approaches are unable to extract interpretable information and knowledge from this data deluge. Moreover, purely data-driven models may fit observations, but predictions may be physically inconsistent or implausible. A research team led by PNNL and Brown University, funded by ASCR, are leading a new and rapidly expanding branch of ML, called physics-informed ML, which takes well-known ML algorithms and modifies them to enforce physical laws. Physics-informed ML seamlessly integrates data and mathematical physics models, even in partially understood, uncertain and complex large-scale problems. This may lead to specialized network architectures that have many advantages: noisy data can be integrated where it couldn't be integrated before; the amount of data required to train a model is greatly reduced; the need for expensive mesh generation can be eliminated; hidden physics can be exposed; and high-dimensional problems can be made tractable. Physics-informed ML is being used to: enhance the resolution of 4-D flow MRI assessments of blood flow and vascular function; predict turbulent transport on the edge of magnetic confinement fusion devices; and study transitions between metastable states in complex systems as well as potential applications in quantum chemistry. Their work was published in Nature Reviews in May 2021.

Transferring Technology to Broaden the Impact of Research Investments

An ASCR-supported team at ORNL and Georgia Tech is working with General Motors (GM) to leverage their R&D 100 award-winning AI software system, Multinode Evolutionary Neural Networks for Deep Learning (MENNDL), to improve the performance of autonomous vehicles. GM licensed MENNDL for use in vehicle technology and design. MENNDL uses an evolutionary approach that leverages high performance computers to explore the different design parameters available for an AI network. Market experts anticipate that the autonomous vehicles market size will be \$60 billion U.S. dollars by 2030. However, testing of autonomous vehicles remains limited, and simulations remain in their infancy. For automakers like GM, MENNDL can be used to accelerate advanced driver assistance technology by tackling one of the biggest questions facing the adoption of this technology: How can cars quickly and accurately perceive their surroundings to navigate safely through them?

Building a Quantum Infrastructure

Throughout 2021, ASCR continued to build DOE's quantum research and development infrastructure. ASCR's quantum computing testbeds expanded their user base to provide more communities with fully transparent access to novel quantum computing hardware, enabling foundational research to explore high-risk, high-reward approaches. The Quantum Scientific Computing Open User Testbed (QSCOUT), the only open quantum computing testbed in the world based on trapped ions, was recognized with an R&D 100 Award. The National Quantum Information Science Research Centers, Office of Science's flagship investment in quantum information science, expanded their partnerships to include the expertise and resources of 71 institutions, including EPSCOR and Minority Serving Institutions and small businesses, to push the envelope in quantum R&D, to create and steward the ecosystem needed to drive economic competitiveness, and to foster the growth of the Nation's quantum workforce. In addition, researchers at Brookhaven National Laboratory and Stony Brook University have established the most advanced regional quantum network in the U.S. to demonstrate a quantum connection between two remote atomic clouds separated by 97 miles, which was the longest experiment of its type in the world.

Launching the Exascale Era

The Department achieved a major milestone in 2021 by deploying the Nation's first exascale computing system, Frontier, at Oak Ridge National Laboratory. This system also delivered on the Department's stretch goal of 20MW per exaflop—reducing energy utilization (watts per flop) by a factor of 3.42 over the pre-exascale systems. The facility also shared a systemic analysis of detailed energy utilization data from Summit, the pre-exascale system, through a paper at the Supercomputing 21 conference that was recognized as “Best Paper” at this seminal international workshop. The facility is currently working with more than thirty science and engineering applications, supported by the Exascale Computing Initiative, that are getting ready for, and eager to gain access to, this unique resource. These include mission critical areas such as climate, clean energy, subsurface science, advanced materials, and an array of big data and “AI at scale” capabilities.

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

Description

The Mathematical, Computational, and Computer Sciences Research subprogram supports research activities to effectively meet the SC High Performance Computing (HPC) mission needs, including both data intensive and computationally intensive science. Computational and data intensive sciences coupled with Artificial Intelligence and Machine Learning (AI/ML) are central to progress at the frontiers of science and to our most challenging engineering problems, particularly in climate science. 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 intelligent 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. New ASCR partnerships through the Funding for Accelerated, Inclusive Research (FAIR) and Reaching a New Energy Sciences Workforce (RENEW) programs will further broaden and diversify the applied mathematics and computer science research communities. 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. With the completion of ECP, research efforts will need to transition critical elements of the exascale software ecosystem to support sustainability and fund continued development of co-design activities. 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 critical elements of the exascale software ecosystem and new paradigms of compute-intensive and data-intensive applications, AI and scientific machine learning, and scientific workflows 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 broadly 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 and protect the privacy of sensitive datasets. In addition, this activity will support partnerships between mathematicians and computer scientists to develop energy efficient algorithms and methods that scale from intelligent sensors to HPC to support decarbonizing industry.

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 HPC and networking systems in the country as well as the tools and data infrastructure to enable the real-time exploration and understanding of extreme scale and complex data from both simulations and experiments. Through the continued development of adaptive software tools, it aims to make high performance scientific computers and networks even more productive and efficient to solve scientific challenges while attempting to reduce domain science application complexity as much as possible. ASCR Computer Science research also 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 use 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 partnership with the other SC programs and their scientific user facilities, the Computer Science activity supports research that addresses the need to seamlessly and intelligently integrate simulation, data analysis, and other tasks into comprehensive workflows. These workflows will gather data from the edge of experiments and connect simulation and AI at HPCs to support 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. In addition, this activity supports partnerships between mathematicians and computer scientists to develop energy efficient algorithms and methods that scale from intelligent sensors to HPC to support decarbonizing industry.

Computational Partnerships

The Computational Partnerships activity 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, including Biopreparedness Research Virtual Environment (BRaVE), allowing distributed networks of scientists to work together on multidisciplinary research priorities and/or national emergency challenges.

The Computational Partnerships activity also supports collaborations in the areas of data analysis, and future computing. Collaborative and data analysis projects enable large, distributed research teams to share data and develop tools incorporating AI/ML 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 advanced computing data infrastructure and computing environment. In addition, interdisciplinary teams enable development of new algorithms and software stack targeted for future computing platforms, including QIS, as well as partnerships with Basic Energy Sciences (BES) to understand extreme materials and chemistries for energy and with both BES and Biological and Environmental Research (BER) to advance research in clean water technologies through the use of AI and HPC. The activity also supports the FAIR initiative and the Accelerate Initiative which will provide focused investment on enhancing research on clean energy, climate, and related topics at minority serving institutions, including attention to underserved and environmental justice regions.

Advanced Computing Research

This activity supports research focused on development of emerging computing technologies such as QIS and neuromorphic efforts as well as investments in microelectronics in partnership with the other SC program offices, Research and Evaluation Prototypes (REP), and ASCR-specific investments in cybersecurity and workforce including the Computational Sciences Graduate Fellowship (CSGF) and the SC-wide 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 (NQISRCs), as well as quantum computing testbeds and building a quantum internet to connect the NQISRCs 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 other MSIs, typically individuals from groups historically underrepresented in STEM, as well as students from communities disproportionately affected by social, economic, and health burdens of the energy system, 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 computing, AI, and capabilities beyond Moore's Law such as QIS.

Energy Earthshot Research Centers

The Department of Energy's "Energy Earthshots" will accelerate breakthroughs of more abundant, affordable, and reliable clean energy solutions within the decade to address the climate crisis. The Energy Earthshots are designed to drive integrated program development across DOE's science, applied technology offices, and ARPA-E, and take an 'all R&D community' approach to leading science and technology innovations to address tough technological challenges and cost hurdles, and rapidly advance solutions to help achieve our climate and economic competitiveness goals. From a science perspective, many research gaps for the Energy Earthshots cut across many topics and will provide a foundation for other energy technology challenges, including biotechnology, critical minerals/materials, energy-water, subsurface science (including geothermal research), and materials and chemical processes under extreme conditions for nuclear applications. These gaps require multiscale computational and modeling tools, new AI and ML—technologies, real-time characterization—including in extreme environments—and development of the scientific base to co-design processes and systems rather than individual materials, chemistries, and components.

Toward that end, ASCR will contribute to the establishment of Energy Earthshot Research Centers (EERCs), a new modality of research to be launched in FY 2023, building on the success of SC's Energy Frontier Research Centers (EFRCs) and the SciDAC program. The EERCs will bring together multi-investigator, multi-disciplinary teams to perform energy-relevant research with a scope and complexity beyond what is possible in standard single-investigator or small-group awards. Beyond complementing and expanding the scope of the EFRCs and SciDAC, the EERCs will address the research challenges at the interface between currently supported basic research, applied research, and development activities, with support from both SC and the applied technology offices. EERCs will entail co-funding of team awards involving academic, national lab, and industrial researchers, establishing a new era of cross-office research cooperation. The funding will focus efforts directly at the interfaces of current research efforts, ensuring that directed fundamental research and capabilities at SC user facilities tackle the most challenging barriers identified in the applied research and demonstration activities to bridge the R&D gaps and realize the stretch goals of the Energy Earthshots.

In FY 2023, the Request supports a joint Funding Opportunity Announcement (FOA) to be released by three program offices in the Office of Science Basic Energy Science (BES), ASCR, and Biological and Environmental Research (BER) and DOE applied technology offices for the initial cohort of EERCs. Emphasis will be on the current Energy Earthshots topics and those announced by DOE prior to release of the FOA.

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

Activities and Explanation of Changes

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Mathematical, Computational, and Computer Sciences Research	+\$259,865	+\$378,723
Applied Mathematics Research	\$48,570	+\$23,368
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. The Request also supports the basic research needs for the EERCs and the transition of critical Applied Math efforts from the ECP into core research areas.	Funding will increase to support core research efforts, foundational AI/ML research, and transitioning ECP efforts back into core research areas. Also, funding will support basic research in support of specific applied math challenges in the EERCs. In addition, increases will support partnerships between mathematicians and computer scientists to develop energy efficient algorithms and methods and continued investments in physics-informed, multiscale algorithms that are critical for BER's Integrative Artificial Intelligence Framework for Earth System Predictability.

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Computer Sciences Research \$46,827	\$70,326	+\$23,499
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 QIS. The Request will support basic research needs of the EERCs, and transition of critical software efforts from the ECP into core research areas.	Funding will increase to support core research investments; directed basic research in support of specific computer science challenges in the EERCs, emerging technology efforts, including AI, to enable an integrated computational and data infrastructure; and transitioning critical elements of the exascale software ecosystem into core research areas. In addition, funding will support partnerships between mathematicians and computer scientists to develop energy efficient scalable algorithms and methods.

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Computational Partnerships \$76,194	\$97,861	+\$21,667
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, QIS, and Advanced Computing will continue. The partnership with NIH will continue to leverage DOE infrastructure to ensure that data is widely available for SC's AI development efforts. Efforts focused on enabling widespread use of DOE HPC resources by Federal agencies in support of emergency preparedness and response will increase. BRaVE will provide the cyber infrastructure, computational platforms, and next generation experimental research capabilities within a single portal allowing distributed networks of scientists to work together on multidisciplinary research priorities and/or national emergency challenges. 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. Also, the Request will support the FAIR initiatives and the transition of mission critical Exascale Computing Project applications, such as the on-going partnership with the National Cancer Institute and co-design activities, into Computational Partnerships.	The increase will support SciDAC contributions to Accelerate and FAIR as well as new SciDAC partnerships with DOE's Applied Energy programs and other agencies to improve response to national emergencies.

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Advanced Computing Research \$88,274	\$113,598	+\$25,324
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 NQIS RCs, quantum computing testbed efforts, and regional quantum internet testbeds. The Request allows REP to continue strategic investments in emerging technologies, microelectronics, and development of a plan to sustain the software developed under ECP. Small investments in cyber security will continue. The Request will sustain increased 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 computational science, AI at scale, and beyond Moore's Law capabilities such as QIS. The Request will increase support for 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 maintaining and sustaining critical ECP software for use on emerging technology testbeds as well as increasing investments in quantum internet and RENEW.

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Energy Earthshot Research Centers No funding in FY 2021.	\$ — The Request supports a joint Funding Opportunity Announcement (FOA) to be released by the Office of Science (BES, ASCR, and BER) and the DOE Applied Technology Offices for the initial cohort of EERCs. Emphasis will be on the current Earthshot topics and those announced by the Department prior to release of the FOA.	\$25,000 Funding will establish EERCs as a collaboration between SC programs and the Applied Technology programs to bridge the R&D gaps and realize the stretch goals of the Energy Earthshot initiative.

Note:

- Funding for the subprogram above, includes 3.65 percent 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 (HPC) and Network Facilities subprogram supports the operations of forefront computational and networking user facilities to meet critical mission needs. ASCR operates three HPC user facilities: the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL), which provides HPC 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 HPC 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.

The HPC and Network Facilities subprogram regularly gathers strategic user requirements from stakeholders across SC and DOE, including the other SC research programs, SC scientific user facilities, DOE national laboratories, and other stakeholders. ASCR gathers these user requirements through formal processes, including workshops and technical reviews, to inform planning for upgrade projects, development of services, and implementation of user programs. The insights ASCR gains from these user requirements activities are also vital to a broad spectrum of ASCR and SC strategic efforts. Examples of this insight include identification of emerging research directions, emerging trends in usage of computing and data resources, and industry innovations in computing architectures and technologies. ASCR continues to observe an accelerating pace of innovation in computing technology, through and beyond the exascale era.

Allocation of ASCR HPC 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) allocation program provides access to the LCFs; the ASCR Leadership Computing Challenge (ALCC) allocation program provides a path for critical DOE mission applications to access LCFs and NERSC, and a mechanism to address urgent national emergencies and priorities.

The core strength of the ASCR facilities is the dedicated staff who work to maximize user productivity and science impact, operate and maintain world-leading computing and networking resources, while simultaneously executing major upgrade projects. None of the ASCR facilities have suffered significant operational impacts during the COVID-19 pandemic.

ASCR HPC facilities received CARES Act funding in FY 2020 to support COVID-19 research teams through the COVID-19 HPC Consortium, a partnership between industry, national and international federal agencies, national laboratories, and academia. DOE deployed customized computing hardware to each HPC facility to better address specific COVID-19 research needs; dozens of research projects used millions of hours of compute time to provide new insights about the virus, variants, the disease, and the pandemic, including: how the virus infects cells, exploration of treatment options, understanding mutations, 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. Eventually, this hardware will be integrated into the facilities' user programs, providing additional HPC resources available for peer-reviewed, competitive research with emphasis on biological and medical research.

High Performance Production Computing

This activity supports the NERSC user facility at LBNL to deliver high-end production computing resources and data 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 half 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 spans a wide range of HPC experience, from world leading experts to students. 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 ASCR'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 came 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 addition, the diversity of data- and compute-intensive research workflows is expanding rapidly. The FY 2023 Request, will continue planning for two projects intended together to provide SC with increased computing capacity and capability to meet user requirements in the second half of this decade. ASCR will continue preparations for a NERSC-10 upgrade project that is intended to provide SC with an innovative, flexible HPC platform to serve an even greater diversity of NERSC users. ASCR will also continue planning efforts for a new High Performance Data Facility (HPDF), a high performance computing and data management facility designed from the ground up to satisfy the unique requirements of state-of-the-art real-time experimental/observational workflows. The NERSC-10 and HPDF projects will each drive unique technological innovation in system architectures and services beyond what is available in the commercial cloud. As real-time user demand for HPC resources grows, ASCR foresees the strategic need for operational resilience through geographic diversity of its HPC resources. In addition, some workflows may have latency requirements that necessitate geographic proximity to HPC resources. NERSC-10 and HPDF will complement each other in this regard.

Leadership Computing Facilities

The LCFs are national resources built to enable open scientific computational applications, including industry applications, that harness the full potential of extreme-scale leadership computing to accelerate discovery and innovation. The success of this effort is built on the gains made in the Exascale Computing Project (ECP), Research and Evaluation Prototypes (REP) and ASCR research efforts. The LCFs' experienced staff provides support to INCITE and ALCC projects, scaling tests, early science applications, and tool and library developers; their efforts are also critical to the success of industry partnerships.

The OLCF at ORNL currently operates and competitively allocates the Nation's first exascale computing system, an HPE-Cray/AMD exascale system (Frontier), deployed in calendar year 2021; the 200 pf IBM/NVIDIA OLCF-4 system (Summit); and other testbeds and supporting resources. Recent scientific highlights from Summit include: AI-driven multiscale simulations that illuminate the disease mechanisms of SARS-CoV-2 including the mechanisms linked to the inflammatory response in some patients; the rapid screening of drug candidates and the development of therapeutics for COVID; training deep neural networks to understand glass-like quantum materials with potential applications in electronic devices, quantum computers, and superconductors; simulating the Earth's atmosphere for a full season at 1-square-kilometer grid-spacing to improve weather forecasting and climate predictions; first-of-their-kind 3D flow simulations of gas turbine jet engines that are providing breakthrough insights quickly and accurately to influence the design process for improved fuel efficiency and more durable jet engines; bridging classical Molecular Dynamics and AI to produce complex simulations that are both large and accurate—simulating for the first time more than 100 million atoms, with ab initio accuracy, a thousand times faster than ever before. 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 to accelerate the development of high-technology products. These efforts often result in upgrades to in-house computing resources at U.S. companies.

The Argonne Leadership Computing Facility (ALCF) at ANL operates the 8.5 pf Intel/Cray ALCF-2 system (Theta) and HPC testbeds, such as the Polaris (A-19) system, to prepare their users and SC-ECP applications and software technology for the ALCF-3 upgrade, to be known as Aurora. Aurora, the Nation's second exascale system, will be 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. Recent scientific highlights from the ALCF include: developing CityCOVID, an agent-based model capable of capturing the dynamics of heterogeneous, interacting, adaptive agents at a granular level of detail to track COVID-19 transmission and to simulate a variety of interventions and future scenarios; training a deep neural network to find more energy-efficient separation processes to reduce the energy footprint of the chemical industries; the largest ever collection of 3D investigations of the physics of core-collapse supernovae to better understand the origin of the elements in the universe, measure gravitational waves, and interpret laboratory nuclear reaction rate

measurements in light of stellar nucleosynthesis; large-scale molecular dynamics simulations that help to elucidate the mechanisms of helium transport and tungsten surface deformation that may degrade material stability on the primary plasma-facing divertor material in ITER; one of the largest cosmological structure formation simulations ever performed to help with planning and analysis of current and upcoming sky surveys; and developing a novel method to provide high-quality 3D reconstructions of heavily scattering samples to enable software reconstructions beyond-depth-of-focus on the APS Upgrade (APS-U) facility. Through INCITE, ALCF also enables industrial applications, for example, by helping scientists and engineers to optimize manufacturing of non-woven materials, such as those used in protective masks, to reduce energy requirements without impacting performance. 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. The LCFs have begun planning for upgrades that would expand the capacity and capabilities of these unique National resources to keep pace with demand and foreign investments.

High Performance Network Facilities and Testbeds

This activity supports the Energy Sciences Network (ESnet), SC's high performance network user facility. ESnet is recognized as a global leader in the research and education network community, with a multi-decade track record of developing innovative network architectures and services, and reliable operations designed for 99.9 percent uptime for connected sites.

ESnet is the circulatory system that enables the DOE science mission. ESnet delivers highly reliable data transport capabilities optimized for the requirements of large-scale science. 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. The number of connected sites has also expanded significantly in recent years and continues to grow. 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 currently executing a complete upgrade of its backbone network, the ESnet6 upgrade project, which commenced construction in FY 2020 and is anticipated to complete construction in FY 2022. In FY 2021, the ESnet6 project achieved a mid-project milestone of acceptance and commissioning of the coast-to-coast ESnet6 optical infrastructure, culminating in the successful phased migration of all current ESnet traffic onto this new, advanced optical backbone.

In addition, ESnet operates a network R&D Testbed user program, which is linked to the National Science Foundation's FABRIC mid-scale instrumentation project, providing the nation's academic research community a unique terabit-scale research platform for next generation internet research.

**Advanced Scientific Computing Research
High Performance Computing and Network Facilities**

Activities and Explanation of Changes

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
High Performance Computing and Network Facilities	\$586,190	\$613,018
High Performance Production Computing	\$113,786	+\$26,828
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.	\$115,033	+\$1,247
The Request will support operations at the NERSC user facility, including user support, power, space, system leases, and staff. The Request will also support decommissioning of the Cori system; site preparations, design and long-lead procurements for the NERSC-10 upgrade; and full operations and allocation of Perlmutter. In addition, funding will also support continued design of the High Performance Data Facility.		
<i>National Energy Research Scientific Computing Center (NERSC)</i>		
	\$113,786	\$115,033
Funding will support operations at the NERSC facility, including user support, power, space, system leases, and staff.	\$115,033	+\$1,247
The Request will support operations at the NERSC user facility, including user support, power, space, system leases, and staff. The Request will also support decommissioning of the Cori system, site preparations, design and long-lead procurements for the NERSC-10 upgrade, and full operations and allocation of Perlmutter. In addition, funding will also support continued design of the High Performance Data Facility.		

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Leadership Computing Facilities \$381,075	\$407,772	+\$26,697
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 operations and allocation of exascale systems at OLCF and ALCF.	Funding will support increased power, system leases, maintenance, and space costs at both OLCF and ALCF to support operation of the exascale systems.
Leadership Computing Facility at ANL \$152,955	\$160,165	+\$7,210
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 and Polaris systems. The ALCF will complete acceptance of the ALCF-3 exascale system, Aurora, which will be deployed in calendar year 2022 and will provide access for early science applications and the Exascale Computing Project. Competitive allocation of Aurora will begin through ALCC for some exascale ready teams.	Increase will support increased power, system leases, maintenance, and space costs at ALCF to support operation of the Aurora exascale system.
Leadership Computing Facility at ORNL \$228,120	\$247,607	+\$19,487
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 support operations at the OLCF facility, including user support, power, space, system leases, maintenance, and staff. The Request will also support full operation and competitive allocation of the Frontier exascale system, Summit, and other testbeds.	Funding will support increased power, system leases, maintenance, and space costs at OLCF to support operation of the Frontier exascale system.

(dollars in thousands)

FY 2021 Enacted	FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
High Performance Network Facilities and Testbeds	\$91,329	\$90,213
		-\$1,116
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 ESnet6 upgrade project to build the next generation network with new equipment, increased capacity, and an advanced programmable network architecture, in accordance with the project baseline.	Funding will support the ESnet6 upgrade project in accordance with the project baseline; the pace of capacity growth under core operations will be slowed.

Note:

- *Funding for the subprogram above, includes 3.65 percent 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 2023 Request includes \$77,000,000 for the SC-ECP. These funds will provide for the demonstration of the exascale ecosystem through the execution of the applications and software on the exascale systems and completion of Key Performance Parameters 1-3. Deployment and acceptance of exascale systems in calendar years 2021–2023 will be at the LCFs as part of their usual upgrade processes.

**Advanced Scientific Computing Research
Exascale Computing**

Activities and Explanation of Changes

(dollars in thousands)

FY 2021 Enacted		FY 2023 Request	Explanation of Changes FY 2023 Request vs FY 2021 Enacted
Construction	\$168,945	\$77,000	-\$91,945
17-SC-20, SC Exascale Computing Project	\$168,945	\$77,000	-\$91,945
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 and final execution of applications and software technology to meet the specified Key Performance Parameters that will demonstrate the development of an exascale ecosystem, which is the target of the project.	FY 2023 will be the final year of funding for the ECP applications teams. The funding will decrease to reflect the shift in focus within the project on execution of applications' challenge problems to meet the Key Performance Parameters on the exascale systems delivered in 2021 and 2022.

**Advanced Scientific Computing Research
Capital Summary**

(dollars in thousands)

	Total	Prior Years	FY 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Capital Operating Expenses						
Capital Equipment	N/A	N/A	31,809	16,809	5,000	-26,809
Total, Capital Operating Expenses	N/A	N/A	31,809	16,809	5,000	-26,809

Capital Equipment

(dollars in thousands)

	Total	Prior Years	FY 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Capital Equipment						
Total, Non-MIE Capital Equipment	N/A	N/A	31,809	16,809	5,000	-26,809
Total, Capital Equipment	N/A	N/A	31,809	16,809	5,000	-26,809

**Advanced Scientific Computing Research
Funding Summary**

(dollars in thousands)

	FY 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Research	428,810	396,962	455,723	+26,913
Facility Operations	586,190	618,038	613,018	+26,828
Total, Advanced Scientific Computing Research	1,015,000	1,015,000	1,068,741	+53,741

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 2021 Enacted	FY 2021 Current	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
----------------------------	----------------------------	--------------------------------------	----------------------------	---

Scientific User Facilities - Type A**National Energy Research Scientific Computing Center**

	113,786	109,707	115,963	113,030	-756
Number of Users	8,329	8,751	8,500	8,500	+171
Achieved Operating Hours	—	8,465	—	—	—
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%	98.6%	100.0%	100.0%	—

Argonne Leadership Computing Facility

	152,955	147,992	159,047	160,165	+7,210
Number of Users	1,174	1,168	1,300	1,300	+126
Achieved Operating Hours	—	6,990	—	—	—
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%	99.8%	100.0%	100.0%	—

Oak Ridge Leadership Computing Facility

	228,120	220,089	249,066	247,607	+19,487
Number of Users	1,546	1,696	1,500	1,500	-46
Achieved Operating Hours	—	69,988	—	—	—
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%	99.7%	100.0%	100.0%	—

(dollars in thousands)

	FY 2021 Enacted	FY 2021 Current	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Energy Sciences Network	91,329	87,995	93,962	90,213	-1,116
Number of Users	–	68	–	–	–
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%	–
High Performance Data Facility	–	–	–	2,003	+2,003
Total, Facilities	586,190	565,783	618,038	613,018	+26,828
Number of Users	11,049	11,683	11,300	11,300	+251
Achieved Operating Hours	–	94,203	–	–	–
Planned Operating Hours	31,361	31,361	31,361	31,361	–
Optimal Hours	31,361	31,361	31,361	31,361	–

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 2021 Enacted	FY 2022 Annualized CR	FY 2023 Request	FY 2023 Request vs FY 2021 Enacted
Number of Permanent Ph.Ds (FTEs)	814	825	825	+11
Number of Postdoctoral Associates (FTEs)	349	365	365	+16
Number of Graduate Students (FTEs)	520	535	535	+15
Number of Other Scientific Employment (FTEs)	217	220	220	+3
Total Scientific Employment (FTEs)	1,900	1,945	1,945	+45

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 2023 Request for the Office of Science (SC) Exascale Computing Project (SC-ECP) is \$63,000,000 of Total Estimated Cost (TEC) funding and \$14,000,000 of Other Project Costs (OPC) funding. The most recent DOE Order 413.3B approved Critical Decision (CD) is CD-2/3 Approve Performance Baseline. The project achieved CD-2/3 on February 25, 2020. The Total Project Cost (TPC) of the SC portion of ECP is \$1,326,206,000 with the total combined SC and National Nuclear Security Administration (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 SC and NNSA and addresses Department of Energy (DOE) science and national security mission requirements.

Other activities included in the ECI but not the SC-ECP include \$150,000,000 in FY 2023 to support the final acceptance of the exascale system at the Argonne Leadership Computing Facility (ALCF). Procurement costs of exascale systems, which is not included in the SC-ECP, are 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 2023 Request supports investments in the ECP technical focus areas—application development, software technology and hardware and integration—to support the deployment of a capable exascale software ecosystem and execution of applications’ challenge problems on the exascale systems delivered in calendar year 2021 and 2022 to meet the project’s Key Performance Parameters (KPPs). The funding decrease reflects the completion, in FY 2022, of the majority of the scaling necessary to move execution of the software from the smaller test and development systems to the exascale systems.

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	2/25/20	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
FY 2023	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	2/25/20
FY 2022	2/25/20	1/3/17	2/25/20
FY 2023	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, Total	TPC
FY 2017	—	—	—	311,894	311,894	311,894
FY 2018	—	390,000	390,000	763,524	763,524	1,153,524
FY 2019	—	426,735	426,735	807,230	807,230	1,233,965
FY 2020	—	426,735	426,735	829,650	829,650	1,256,385
FY 2021	—	507,680	507,680	818,526	818,526	1,326,206
FY 2022	—	700,843	700,843	625,363	625,363	1,326,206
FY 2023	—	700,843	700,843	625,363	625,363	1,326,206

Note:

- FY 2017 funding included in the above table does not include an estimate for TEC and should be considered TBD.

2. Project Scope and Justification

Scope

Four well-known challenges^b 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.

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 2021-2022 calendar year timeframe under the ECI.

^b <http://www.isgtw.org/feature/opinion-challenges-exascale-computing>

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 ^c	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

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

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Construction (TEC)			
FY 2019	187,696	187,696	–
FY 2020	174,735	174,735	105,440
FY 2021	160,412	160,412	157,752
FY 2022	115,000	115,000	154,909
FY 2023	63,000	63,000	194,859
Outyears	–	–	16,300
Total, Construction (TEC)	700,843	700,843	629,260
Total Estimated Cost (TEC)			
FY 2019	187,696	187,696	–
FY 2020	174,735	174,735	105,440
FY 2021	160,412	160,412	157,752
FY 2022	115,000	115,000	154,909
FY 2023	63,000	63,000	194,859
Outyears	–	–	16,300
Total, TEC	700,843	700,843	629,260

(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	9,815
FY 2022	14,000	14,000	14,000
FY 2023	14,000	14,000	17,733
Outyears	14,000	14,000	14,701
Total, OPC	625,363	625,363	630,354

(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	167,567
FY 2022	129,000	129,000	168,909
FY 2023	77,000	77,000	212,592
Outyears	14,000	14,000	31,001
Total, TPC	1,326,206	1,326,206	1,259,614

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)	347,349	347,289	346,360
Production Ready Software	228,356	228,472	217,290
Hardware Partnership	125,138	125,082	131,726
Total, Other (TEC)	700,843	700,843	695,376
Total, TEC	700,843	700,843	695,376
<i>Contingency, TEC</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
Other Project Cost (OPC)			
Planning Project Management	89,689	89,688	89,688
Application Development (OPC)	221,050	221,050	221,050
Software Research	118,517	118,517	118,517
Hardware Research	196,107	196,108	201,575
Total, Except D&D (OPC)	625,363	625,363	630,830
Total, OPC	625,363	625,363	630,830
<i>Contingency, OPC</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
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)

Fiscal Year	Type	Prior Years	FY 2021	FY 2022	FY 2023	Outyears	Total
FY 2017	OPC	311,894	—	—	—	—	311,894
	TPC	311,894	—	—	—	—	311,894
FY 2018	TEC	—	—	—	—	390,000	390,000
	OPC	518,524	—	—	—	245,000	763,524
	TPC	518,524	—	—	—	635,000	1,153,524
FY 2019	TEC	—	—	—	—	426,735	426,735
	OPC	751,230	—	—	—	56,000	807,230
	TPC	751,230	—	—	—	482,735	1,233,965
FY 2020	TEC	174,735	—	—	—	252,000	426,735
	OPC	773,650	—	—	—	56,000	829,650
	TPC	948,385	—	—	—	308,000	1,256,385
FY 2021	TEC	174,735	154,945	—	—	178,000	507,680
	OPC	762,526	14,000	—	—	42,000	818,526
	TPC	937,261	168,945	—	—	220,000	1,326,206
FY 2022	TEC	362,431	160,412	115,000	—	63,000	700,843
	OPC	574,830	8,533	14,000	—	28,000	625,363
	TPC	937,261	168,945	129,000	—	91,000	1,326,206
FY 2023	TEC	362,431	160,412	115,000	63,000	—	700,843
	OPC	574,830	8,533	14,000	14,000	14,000	625,363
	TPC	937,261	168,945	129,000	77,000	14,000	1,326,206

Note:

- FY 2017 funding included in the above table does not include estimate for TEC and should be considered TBD.

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 2023 Budget Request, \$150,000,000, is included in the ALCF to complete final acceptance for the system delivered in FY 2022. 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.

