

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. ASCR's partnerships within the Office of Science (SC) and with the applied technology offices, other agencies, and industry are essential to these efforts. 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. 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, DOE and the Nation. ASCR also deploys and operates world-class, open access HPC 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 them. The benefits of U.S. computational leadership have been enormous 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. There is recognition that the nation that leads in Artificial Intelligence (AI) and in the integration of the computing and data ecosystem will lead the world in developing innovative clean energy technologies, medicines, industries, 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. The next generation of breakthroughs in computational 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 talents across SC, universities, national labs, and the private sector in concerted efforts to enable the U.S. to lead the world into the quantum future.

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 (ECP); 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 2024 Request

The FY 2024 Request of \$1,126.0 million for ASCR is well-aligned with the Administration and the Department priorities to catalyze research and innovation in critical and emerging technologies (AI, QIS, HPC, advanced communications technologies, and microelectronics), Energy Earthshots, and pandemic readiness and prevention.

Research

- The Request prioritizes foundational research in Applied Mathematics and Computer Science to ensure critical technologies from the ECP are maintained and improved so the full potential of the exascale systems can be realized. Investments will emphasize foundational research to address the combined challenges of increasingly heterogeneous and reconfigurable architectures and the changing ways in which HPC systems are used, advancing high efficiency computing, incorporating AI into simulations and data intensive applications while minimizing bias, and increasing connectivity and integration across distributed resources. The Request supports new Microelectronics Science Research Centers as authorized under the Micro Act, focusing on a multi-disciplinary co-design innovation ecosystem in which materials, chemistries, devices, systems, architectures, algorithms, and software are developed in a closely integrated fashion. Support increases for the new Energy Earthshot Research Centers (EERCs). The Computational Partnerships activity will continue to partner with other programs to infuse the latest developments in applied math and computer science into strategic applications, including areas such as accelerating the development of emerging technologies, tackling climate change, pandemic preparedness, cancer research, and efforts with DOE's applied technology programs. The Request continues support for research in states and territories with historically lower levels of Federal academic research funding through the Established Program to Stimulate Competitive Research (EPSCoR) program. ASCR funding opportunities will prioritize sustaining ECP teams, software, technologies, and applications as ECP concludes.
- The Request provides robust support for Advanced Computing Research's quantum investments and partnerships in the National Quantum Information Sciences Research Centers (NQISRCs), quantum internet, and testbeds. ASCR's regional quantum testbeds and user programs, which provide U.S. researchers with access to unique and commercial quantum computing and networking resources, and basic research in quantum information will continue to provide national leadership in quantum in coordination with relevant agencies. Support for the Computational Sciences Graduate Fellowship (CSGF) is increased. The Reaching a New Energy Sciences Workforce (RENEW) initiative expands targeted efforts, including a RENEW graduate fellowship, to broaden participation and advance justice, equity, diversity, and inclusion in SC-sponsored research.

Facility Operations

- FY 2024 request supports operations and competitive allocation of the Nation's exascale computing systems at the Oak Ridge Leadership Computing Facility (OLCF), a system called Frontier deployed in calendar year 2021, and Argonne Leadership Computing Facility (ALCF), a system called Aurora that was deployed in calendar year 2022. Funding also supports operations at the National Energy Research Scientific Computing Center (NERSC) and the Energy Sciences Network (ESnet). The Request supports advanced computing and AI testbeds at the facilities with competitive, merit reviewed, open access for researchers. ASCR facilities will deploy ECP software and technologies and will prioritize sustaining ECP software and technology critical to HPC operations and users as ECP concludes. In addition, the request supports continued planning for upgrades for the Leadership Computing Facilities and the NERSC-10 upgrade, including site preparations, long lead procurements, and vendor R&D partnerships, to address rising demand for computing and U.S. competitiveness in advanced computing and computational science and a new High Performance Data Facility to provide crucial resources to SC programs to attack fundamental problems in science and engineering.
- ASCR will continue planning for SC's Integrated Research Infrastructure for state-of-the-art real-time experimental/observational workflows and to 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 Leadership Computing Facilities (LCFs).

Projects

- The ASCR FY 2024 Request includes \$14.0 million for SC's contribution to ECP project including Critical Decision (CD)-4 activities and lessons learned, focused on delivering the Key Performance Parameters (KPPs), open source publishing of ECP software and technologies, and working with vendors and other partners to transfer ECP technologies.
- The FY 2024 Request initiates the HPDF to support the next step in design and planning, advancing from pre-conceptual design to conceptual design, site selection, and alternative selection, and potentially, commencement of site preparation.

**Advanced Scientific Computing Research
Funding**

(dollars in thousands)

	FY 2022 Enacted	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Advanced Scientific Computing Research				
Applied Mathematics Research	50,858	61,035	76,188	+15,153
Computer Sciences Research	49,963	60,667	86,017	+25,350
Computational Partnerships	79,456	95,875	87,600	-8,275
Advanced Computing Research	105,723	108,920	149,848	+40,928
Energy Earthshot Research Centers	–	12,500	12,500	–
Total, Mathematical, Computational, and Computer Sciences Research	286,000	338,997	412,153	+73,156
High Performance Production Computing	120,000	132,003	142,000	+9,997
Leadership Computing Facilities	410,000	430,000	466,607	+36,607
High Performance Network Facilities and Testbeds	90,000	90,000	90,213	+213
Total, High Performance Computing and Network Facilities	620,000	652,003	698,820	+46,817
17-SC-20, SC Exascale Computing Project	129,000	77,000	14,000	-63,000
Subtotal, Advanced Scientific Computing Research	1,035,000	1,068,000	1,124,973	+56,973
Construction				
24-SC-20 High Performance Data Facility	–	–	1,000	+1,000
Subtotal, Construction	–	–	1,000	+1,000
Total, Advanced Scientific Computing Research	1,035,000	1,068,000	1,125,973	+57,973

SBIR/STTR funding:

- FY 2022 Enacted: SBIR \$28,194,000 and STTR \$3,965,000
- FY 2023 Enacted: SBIR \$10,112,000 and STTR \$1,422,000
- FY 2024 Request: SBIR \$12,093,000 and STTR \$1,701,000

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

(dollars in thousands)

FY 2024 Request vs FY 2023 Enacted

Mathematical, Computational, and Computer Sciences Research

The Computer Science and Applied Mathematics activities will: continue foundational and long-term basic research efforts that explore and prepare for emerging technologies; transition critical technologies from the ECP into core research efforts; develop new scalable energy efficient algorithms and software; and address the challenges of data intensive science and emerging computing technologies such as quantum information science, as well as the development of critical tools, including AI, to enable an integrated computational and data infrastructure. Efforts to address specific basic research challenges for the EERCs are increased. Computational Partnerships continues to support partnerships across SC and interagency partnerships such as cancer and biopreparedness. The EERCs will expand basic research to address new challenges. The Advanced Computing Research activity will support the microelectronics research centers, NQISRCs, quantum testbeds, and regional quantum networking testbeds, in close coordination with the other SC programs. Support for CSGF will increase. RENEW expands its targeted efforts to increase participation and retention of underrepresented groups in SC research activities, including through a RENEW graduate fellowship. Funding opportunities will prioritize sustaining ECP teams, software, technologies, and applications as ECP concludes, as researchers shift from project funding to program funding.

+\$73,156

High Performance Computing and Network Facilities

The OLCF and ALCF will provide full operations and competitive allocation of the nation’s exascale computing systems, Frontier and Aurora. Both facilities will deploy and sustain ECP software and technologies critical to operations and will provide testbed resources to explore emerging technologies. In addition, funding supports operation of the 125 petaflop NERSC-9 Perlmutter system. Planning, site preparations and project efforts for NERSC-10. Funding supports operations of all facilities—including power and cooling, equipment, staffing, testbeds, lease payments, user programs and outreach. This activity will continue implementation of SC’s Integrated Research Infrastructure, including planning for a High Performance Data Facility (HPDF).

+\$46,817

Exascale Computing

The FY 2024 Request supports activities to close out and document DOE’s ECP including CD-4 and lessons learned.

-\$63,000

Construction

The FY 2024 Request initiates conceptual design activities for the HPDF project.

+\$1,000

Total, Advanced Scientific Computing Research

+\$57,973

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. ASCR continues to have a strong partnership with National Nuclear Security Administration (NNSA) for achieving the Department's goals for exascale computing. There are growing areas of collaboration across DOE and with other agencies in the area of data-intensive science, AI, and readying applications for exascale. Through the Networking and Information Technology R&D Subcommittee of the National Science and Technology Council (NSTC) Committee on Technology, ASCR also coordinates with programs across the Federal Government. Future Advanced Computing, Scientific Data, Large Scale Networking, AI, High End Computing, and QIS are coordinated with other agencies through the NSTC. In FY 2024, cross-agency interactions and collaborations will continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Delivering on the Exascale Computing Project

ECP teams have successfully run mission critical science and engineering applications on the Nation's first exascale computer, Frontier, at the OLCF. Using Frontier, the ECP WarpX team won the Association for Computing Machinery's (ACM's) 2022 Gordon Bell Prize at the 2022 International Conference for High Performance Computing, Networking, Storage, and Analysis, or SC22, in Dallas, Texas, in November for their outstanding achievement in high-performance computing. WarpX, developed by researchers from Lawrence Berkeley National Laboratory (LBNL), Lawrence Livermore National Laboratory (LLNL), and the French Alternative Energies and Atomic Energy Commission (CEA), is the first mesh-refined, particle-in-cell code for kinetic plasma simulations that is optimized for parallel computing. The team produce a 3D simulation at scale of their own novel concept: a combined plasma particle injector and accelerator that focuses a high-power femtosecond (1 quadrillionth of a second) laser onto a hybrid solid/gas target. Their code ran more than 500 times faster than their original code, Warp, after six years of modernizing and optimizing as an ECP application. In addition, ORNL researchers designed a machine-learning software stack that predicts how strongly a given drug molecule will bind to a pathogen as well as the 3D structure of how it will attach to the target. These vital pieces of information can greatly shorten the usual trial-and-error process of lab experimentation to find viable drug candidates, especially for novel viruses with unknown structures. The software, TwoFold, was a finalist for the 2022 Gordon Bell Special Prize for HPC-Based COVID-19 Research.

Broadening Participation in our Exascale Future

Across the DOE labs and the Nation, demand greatly outpaces supply for highly skilled workers needed to realize the promise of Exascale computing. Growing this workforce requires engaging with more talented people at key stages in their career. The DOE ECP is helping to build a more vibrant, diverse, and inclusive workforce in HPC through a multipronged initiative that is reaching out to a diverse group of talented people from underrepresented communities, including people who are Black or African American, Hispanic/Latinx, American Indian, Alaska Native, Native Hawaiian, Pacific Islanders, women, persons with disabilities, first-generation scholars, and people from smaller colleges and universities. During the summer 2022, 16 faculty and 45 students representing 32 institutions (with 82 percent representing at least one element of diversity) participated in Sustainable Research Pathways for HPC. This effort, initiated by the ECP and transitioned to RENEW, is designed to connect students and faculty from underrepresented groups with DOE lab scientists to encourage lasting collaborations and jump start careers. Participants worked with ECP teams across nine DOE labs to build software technologies that power HPC discoveries and to develop advanced simulation capabilities across mission areas, with results presented by the students at the ECP Annual Meeting.

Speeding Up Quantum Chemistry for Climate-Change Resilience

Understanding how plants respond to drought conditions is critical to adapting to climate change in many parts of the world, and work using quantum-chemical modeling to understand how calcium is transported across cell membranes hopes to shed light on this important topic. A novel graph algorithm developed by a computer-science research project, led by Purdue University, in collaboration with researchers at Pacific Northwest National Laboratory (PNNL) working on the ExaGraph and NWChemEX efforts within the Exascale Computing Project, made a key step in the modeling of a relevant protein, four times faster and enabled the NWChemEX application to scale from using only 4,000 processors to using 14,000 processors on the Summit supercomputer at Oak Ridge National Laboratory. The algorithm, published in 2001 in the Proceedings of Society for Industrial and Applied Mathematics Applied Computational Discrete Algorithms, which computes

an approximate submodular b-matching, accomplishes this feat in NWChemEX by improving the method through which computational tasks are assigned to processors.

Moving Toward an Integrated Research Infrastructure - ESnet-6 enables Distributed Grid Simulation

In July 2022, the ESnet upgrade project reached successful completion ahead of schedule and under budget with enhanced capability and capacity that provided the fast, predictable timing and programmable flexibility essential for reliable exchange of command-and-control information between National Renewable Energy Laboratory's (NREL) Advanced Research on Integrated Energy Systems (ARIES) platform and PNNL during a successful multi-laboratory demonstration of a complex energy system simulation. The FY 2022 multi-laboratory demonstration showed that advanced control systems in the Cordova microgrid could allow it to maintain power to critical resources like the hospital and the airport during an extreme weather event with loss of some of its hydropower resources. During this tightly coupled experiment, NREL simulated the Cordova microgrid while PNNL simulated the advanced control systems and ESnet delivered consistent, low-latency performance critical to the success of this complex, distributed experiments. The team is looking to next connect six national laboratories for experiments that use and model millions of interconnected devices. Such an integration could better identify solutions for city- and region-level problems, with experiments on outages caused by extreme weather events or a cyberattack and to help energy systems managers adopt new technologies that rely on interconnected devices.

Advancing Reliable Fusion Reactor Design and Operation for Sustainable Clean Energy

Nuclear fusion could provide a sustainable supply of clean energy. A critical piece of fusion reactor design is the divertor, which is the only part of the reactor where the extremely hot plasma comes into direct contact with the reactor vessel. The divertor for the multi-billion dollar international fusion experiment ITER will face heat fluxes of 10-20 MW per square meter - ten times higher than the heat load on a spacecraft re-entering Earth's atmosphere. The heat-load width of the divertor is a key design parameter which will sustain exposure to repeated hot exhaust particles. To demonstrate that the tungsten components can withstand the demanding thermal conditions of the ITER machine, researchers supported by the ECP and SciDAC, have integrated high fidelity simulations together with supervised machine learning from experimental data to make more precise predictions. The team used half of the OLCF's Summit supercomputer for two days to complete new simulations that showed that at full power ITER's divertor heat-load width would be more than six times wider than was expected. This would allow ITER to make much faster scientific progress at lower cost. The team produced a more complete prediction of ITER's divertor heat-load width by including magnetic fluctuation and electrostatic turbulence with machine learning parameters to validate the increase predicted for ITER's heat-load width at full power and produced the same results as previous experiments and simulations for existing tokamaks. The results were published in *Physics of Plasmas*.

Mitigating Errors to Advance Quantum Computing

Quantum computers have been improving rapidly over the past several years but they still exhibit error rates far higher than conventional computers. In quantum computers, noise from magnetic fields, changes in temperature, and other sources, leads to the accumulation of errors in complex quantum simulations. This limits their usefulness. To move quantum computing forward, researchers at LBNL developed a new approach to quantum error mitigation that could help make quantum computing's theoretical potential a reality: noise estimation circuits. When combined with three other error mitigation techniques, noise elimination circuits obtained reliable results for dynamic simulations of materials. The novel error mitigation approach will allow researchers to run longer, more realistic simulations and still obtain reliable results. This will broaden the potential impact of upcoming quantum computers on scientific discovery in a huge range of fields, from clean energy to artificial intelligence.

Integrating Powerful New Tools to Accelerate Discovery Science

Metals crack, neurons misfire, viruses mutate—all at scales of size and time we can barely fathom, let alone study. To understand how processes work or fail in the natural or mechanical world, requires the ability to probe more deeply and expose layers of detail never before observed. As our tools advance, the data generated quickly outpaces our ability to keep pace. A cross-cutting team from Argonne National Laboratory (ANL) has realized automated end-to-end analysis of data from the Advanced Photon Source (APS). The integration of APS instruments with ALCF computation and storage resources, leverages data acquisition and management software developed at the APS, the DOE-funded Globus toolkit, and ALCF knowhow to enable automated capture of data at the APS, transfer of data to the ALCF where it is analyzed, and sharing results with the scientific team—all in near real-time without human intervention. This achievement, applied at the upgraded APS in collaboration with the new ALCF Aurora supercomputer, will enable leveraging advanced Artificial

Intelligence/Machine Learning (AI/ML) methods, unlock new scientific opportunities, and enable scientific exploration at speeds and scales previously inaccessible to researchers.

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 HPC and computational science 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 for the Energy Earthshots and 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 needs. 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, which is much longer than the lifespan of the average HPC system. Research efforts must therefore anticipate changes in hardware and rapidly developing capabilities such as AI and QIS, as well as science needs over the long term. ASCR's partnerships with vendors and discipline sciences are essential to these efforts. Within available resources, ASCR will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes. In part through continued funding for the EPSCoR, RENEW and Funding for Accelerated, Inclusive Research (FAIR) initiatives, ASCR will build stronger programs with underrepresented institutions and regions, including investing in a more diverse and inclusive workforce. 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, while minimizing bias;
- 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 advance the Department's energy goals.

Computer Science 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 understanding gaps and future opportunities for the design of future computing systems that maintains U.S. leadership in high-performance and data-intensive computing and developing and evolving the sophisticated software required for these systems, including basic research for diverse computing architectures such as 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 advance the Department's energy goals.

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, simulation, large-scale data analysis, and AI and scientific machine learning in areas of strategic importance to SC, DOE, and the Nation.

The Computational Partnerships activity also supports collaborations to 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 an integrated research infrastructure and computing environment. The activity also supports the FAIR and Accelerate initiatives, which provide focused investment on enhancing research on clean energy, climate, and related topics, including attention to underserved and environmental justice regions and Historically Black Colleges and Universities (HBCUs) and minority serving institutions (MSIs), as well as Biopreparedness Research Virtual Environment (BRaVE) that advances collaborative research for epidemiology frameworks, computational modeling, and data management/integration in support of national biopreparedness and emergency challenges. BRaVE also supports the incorporation of AI/ML and HPC in cancer research in partnership with the National Cancer Institute.

Additionally, this activity provides support for the DOE EPSCoR that funds research in states and territories with historically lower levels of Federal academic research funding. In FY 2024, the EPSCoR program will focus on EPSCoR State-National Laboratory Partnership awards to promote single principal investigator (PI) and small group interactions with the unique capabilities of the DOE national laboratory system and continued support of early career awards.

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 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 enabling DOE researchers to better understand the challenges and capabilities of emerging

technologies. In addition to REP, this activity supports ASCR's investments in the NQISRCs, as well as quantum computing testbeds and quantum internet testbeds.

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 participates 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. Science and Technology (S&T) ecosystem. This includes HBCUs and MSIs, typically individuals from groups historically underrepresented in Science, Technology, Engineering and Math (STEM), as well as students from communities disproportionately affected by social, economic, and health burdens of the energy system. 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.

This activity also provides support for the DOE EPSCoR that funds research in states and territories with historically lower levels of Federal academic research funding. In FY 2024, the EPSCoR program will focus on EPSCoR State-National Laboratory Partnership awards to promote single PI and small group interactions with the unique capabilities of the DOE national laboratory system and continued support of early career awards.

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 Advanced Research Projects Agency-Energy (ARPA-E), and take an 'all research and development (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 continue to partner with SC's Basic Energy Sciences (BES) and Biological and Environmental Research (BER) programs in support of the EERCs, a new modality of research launched in FY 2023, building on the success of SC's Energy Frontier Research Centers (EFRCs) and the SciDAC program. The EERCs 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 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 entail co-funding of team awards involving academic, national laboratories, and industrial researchers, establishing a new era of cross-office research cooperation. Joint funding focuses efforts directly at the interfaces of current research, 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.

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

Activities and Explanation of Changes

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Mathematical, Computational, and Computer Sciences Research	+\$338,997	+\$412,153
Applied Mathematics Research	\$61,035	+\$15,153
Funding continues 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. Funding also supports the basic research needs for the EERCs and the transition of critical Applied Math efforts from the ECP into core research areas.	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 will continue partnerships between mathematicians and computer scientists to develop energy efficient algorithms and methods and investments in physics-informed, multiscale algorithms. The Request also will increase support for the basic research needs for several EERCs. Across this activity, within available resources, efforts will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes.	Funding will increase support for basic research that addresses specific cross-cutting applied math challenges that support several EERCs and transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes.

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Computer Science Research	\$60,667	\$86,017 +\$25,350
<p>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 also continues long-term basic research efforts that explore and prepare for emerging technologies, such as quantum networking, specialized and heterogeneous hardware and accelerators, and QIS. Funding supports basic research needs of the EERCs, and transition of critical software efforts from the ECP into core research areas.</p>	<p>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 computing and networking, and other specialized and heterogeneous hardware and accelerators. Interdisciplinary quantum computing research programs previously under Computational Partnerships are moved to Computer Science to create better synergies. In addition, funding will support partnerships between mathematicians and computer scientists to develop energy efficient scalable algorithms and methods. The Request will increase support for the basic research cross-cutting needs of the EERCs. Across this activity, within available resources, efforts will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes.</p>	<p>Funding will increase to support directed basic research in support of specific cross-cutting computer science challenges in several EERCs and transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes. In addition, funding will support interdisciplinary quantum computing research programs that were previously under Computational Partnerships.</p>

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Computational Partnerships \$95,875	\$87,600	-\$8,275
<p>Funding continues support for the SciDAC Institutes and partnerships with SC and DOE applications. Partnerships on scientific data, AI, QIS, and Advanced Computing continues. The partnership with NIH continues 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 are increased. BRaVE provides 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, as well as transitioning ECP capabilities, such as the on-going partnership with the National Cancer Institute. Also, the funding supports the FAIR initiative with new EPSCoR awards fostering partnerships with national laboratories to leverage unique capabilities of the DOE national laboratory system.</p>	<p>The Request will continue support for the SciDAC Institutes and partnerships with SC and DOE applications. Support for Advanced Computing will continue. Efforts focused on enabling widespread use of DOE HPC resources by Federal agencies in support of emergency preparedness and response will continue. BraVE will provide the cyber infrastructure, computational platforms, and next generation experimental research capabilities to allow networks of scientists to work together on multidisciplinary research priorities and/or national emergency challenges, such as the on-going partnership with the National Cancer Institute. Also, the Request will support the FAIR and Accelerate initiatives, including EPSCoR State-National Laboratory Partnership awards. Interdisciplinary quantum computing research programs are moved to Computer Science to create better synergies. Across this activity, within available resources, ASCR will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes.</p>	<p>Decrease reflects shift of interdisciplinary quantum computing research to Computer Science. Continued support for research in EPSCoR jurisdictions.</p>

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Advanced Computing Research	\$108,920	\$149,848 +\$40,928
<p>Funding continues to support the NQISRCs, quantum computing testbed efforts, and regional quantum internet testbeds. Funds allow REP to continue strategic investments in emerging technologies, microelectronics, and development of a plan to sustain the software developed under ECP. Small investments in cybersecurity continue. Funding sustains increased support for the CSGF fellowship, in partnership with NNSA, supporting increased tuition costs, in order 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. Funding increases support for the RENEW initiative providing undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem, including EPSCoR institutions and students, thus expanding the pipeline for ASCR research and facilities workforce needs.</p>	<p>The Request will continue to support the NQISRCs, quantum computing testbed efforts, and regional quantum internet testbeds. The Request allows REP to increase strategic investments in emerging technologies, microelectronics, and development of a plan to sustain the software developed under ECP. Small investments in cybersecurity will continue. The Request will increase support for the CSGF fellowship, in partnership with NNSA, to support increased tuition costs, an increased stipend, and 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 also 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. Funding will support EPSCoR State-National Laboratory Partnership awards and early career awards. Within available resources, this activity will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes. New Microelectronics Science Research Centers are established, as authorized under the Micro Act.</p>	<p>The Request supports increases for new microelectronics research centers, CSGF, and RENEW. The Request will prioritize transitioning ECP researchers, software, and technologies into core research efforts and DOE priority research areas as ECP concludes. Continued support for research in EPSCoR jurisdictions.</p>

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Energy Earthshot Research Centers	\$12,500	\$12,500
Funding 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 is on the current Earthshot topics and those announced by the Department prior to release of the FOA.	The Request supports joint efforts between Office of Science program (BES, ASCR, and BER) with strong coordination the DOE Applied Technology Offices for the EERCs. Emphasis will be on the current Earthshot topics as well as those announced by the Department prior to release of the FOA.	EERC efforts will inform foundational research investments in applied mathematics and computer science that address the longer-term challenges of the Energy Earthshots.

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, excluding facility operations.*

Advanced Scientific Computing Research High Performance Computing and Network Facilities

Description

The 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 NERSC at LBNL, which provides HPC resources and large-scale storage to a broad range of SC researchers; and the two LCFs at ORNL and 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 the 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.

In FY 2024, the ASCR facilities will continue planning and begin implementation to advance DOE's Integrated Research Infrastructure (IRI) so that researchers can seamlessly and securely meld DOE's unique data, user facilities, and computing resources to accelerate discovery and innovation.

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 10,000 computational scientists conducting over 1,000 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 125 pf HPE/AMD/NVIDIA NERSC-9 system (Perlmutter), which 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 four to five years.

In addition, the diversity of data- and compute-intensive research workflows is expanding rapidly. The FY 2024 Request supports the NERSC-10 upgrade project, which 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 DOE's IRI to satisfy the unique requirements of state-of-the-art real-time experimental/observational workflows and data-integration intensive workflows across the SC user facilities. As demand for HPC resources grows and diversifies, ASCR foresees the strategic need for operational resilience and software portability across its HPC resources.

In FY 2024, the High Performance Production Computing activity will continue planning for the HPDF. Upgrades across the SC user facilities will produce unprecedented increases in data generation. SC programs and their Scientific User Facilities have proposed developing AI/ML techniques to steer experiment and facilities, as well as to speed scientific discovery by automating and streamlining interpretation of datasets. Interaction with experiments in real time requires a service type that existing facilities cannot provide such as the ability to guarantee a computing resource and quality of service throughout an experiment. Effective use of AI/ML also requires the confluence of large well-curated datasets and the compute resources to perform net training activities. Currently, most analyses of experimental and simulation data are done after the experiment or simulation has run. Controlling experiments with AI requires low-latency analysis and inference using high-volume, high-velocity data sets in real time. The proposed HPDF, serving as the foundation for the IRI, will provide a crucial resource to SC programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources. The facility will be designed to dynamically configure computation, network resources, and storage to access data at rest or in motion, supporting the use of well-curated datasets as well as near real-time analysis on streamed data directly from experiments or instruments.

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 ECP, Research and Evaluation Prototypes (REP) and ASCR research efforts. The LCFs' experienced staff deploy cutting edge technologies and provide support to users including ECP teams, scaling tests, early science applications, and tool and library developers; their efforts are also critical to the success of industry and interagency 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: efforts, in partnership with the National Cancer Institute, to train privacy-preserving transformer models for clinical natural language processing of cancer records from the Surveillance, Epidemiology, and End Results (SEER) Program; prediction of synergistic drug combinations for treatment of COVID-19; quantum simulations of photosystem II and cuprate superconductivity; extreme-scale simulations for advanced seismic ground motion and hazard modeling; the world's first seasonal timescale global simulation of the Earth's atmosphere with 1 kilometer average grid spacing; exascale simulation and deep learning models for energetic particles in burning plasmas; design of novel titanium based alloys for additive manufacturing using HPC-aided large-scale phase field simulations; and high-fidelity simulations of turbulent aeroacoustics enabling sustainable aviation. OLCF staff shares its expertise with industry to broaden the benefits of exascale 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 prompt U.S. companies to expand their own HPC resources.

The ALCF at ANL operates and competitively allocates the Nation's second exascale system, an Intel/HPE-Cray system (Aurora) deployed in calendar year 2022. ALCF also operates and allocates a 44 PF HPE/AMD/NVIDIA testbed (Polaris) to prepare users for Aurora and to support large-scale data analytics and machine learning. The ALCF also operates a versatile AI testbed consisting of five systems, leading the investigation and sharing of novel AI accelerators and facilitating publications, and supported a 2021 COVID-19 Gordon Bell finalist with these pre-exascale systems. Recent scientific highlights from the ALCF include: the first tests of novel chiral nucleon-nucleon potentials consistent with three-nucleon interactions in a critical step in developing a first-principles description of nuclear structure; identification of an electrolyte that can be used to protect lithium-ion batteries from water damage; new geophysical dynamic rupture models to perform more accurate seismic hazard analysis; and a digital twin of the city of Chattanooga, Tennessee, to find more effective ways to improve energy efficiency for the city's buildings. The ALCF and OLCF systems are architecturally distinct, consistent with DOE's strategy to manage enterprise risk, foster diverse capabilities that provide the Nation's HPC user community with the most effective resources, and expand U.S. competitiveness.

The demand for 2022 INCITE allocations at the LCFs outpaced the available resources by a factor of three, and this is expected to increase with the availability of Aurora, and its unique capabilities. Demand for 2021–2022 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.

In FY 2024, the LCFs will continue planning for future upgrades, cultivate vendor partnerships to spur innovation of strategic value and drive U.S. competitiveness, and contribute to planning DOE’s integrated research infrastructure.

High Performance Network Facilities and Testbeds

This activity supports ESnet, SC’s high performance network user facility, providing world-leading wide-area network access for all of DOE. ESnet is widely 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 recently completed a major upgrade of its backbone network, the ESnet6 project, which commenced construction in FY 2020, launching a new era of automation, and programmability that will provide DOE science with higher performance and drive greater optimization of network resources.

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

In FY 2024, ESnet will leverage the unique attributes of ESnet6 to develop advanced services to support DOE priority R&D thrusts, DOE’s IRI, and cybersecurity.

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

Activities and Explanation of Changes

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
High Performance Computing and Network Facilities	\$652,003	\$698,820
		+\$46,817
High Performance Production Computing	\$132,003	\$142,000
		+\$9,997
Funding supports operations at the NERSC user facility, including user support, power, space, system leases, and staff. Funding also supports 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 supports continued design of the HPDF.	The Request will support operations at the NERSC user facility, including user support, power, space, system leases, and staff. NERSC will deploy the exascale computing software and will prioritize sustaining ECP software and technologies critical to HPC operations and users as ECP concludes. The Request will also support activities such as site preparations, design and procurements for the NERSC-10 upgrade. In addition, funding will support continued design of DOE's IRI.	The funding increase will support site preparations, design and procurements for the NERSC-10 upgrade, continued planning for DOE's IRI, and sustaining ECP software and technologies critical to HCP operations and users.
<i>National Energy Research Scientific Computing Center (NERSC)</i>	<i>\$130,000</i>	<i>\$135,000</i>
		<i>+\$5,000</i>
Funding supports operations at the NERSC user facility, including user support, power, space, system leases, and staff. Funding supports 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 supports continued design of the HPDF.	The Request will support operations at the NERSC user facility, including user support, power, space, system leases, and staff. NERSC will deploy the exascale computing software and will prioritize sustaining ECP software and technologies critical to HPC operations and users as ECP concludes. The Request will also support activities such as site preparations, design and procurements for the NERSC-10 upgrade, and full operations and allocation of Perlmutter. In addition, funding will also support continued design of DOE's IRI.	Funding will support site preparations, design and long-lead procurement for the NERSC-10 upgrade, continued planning for DOE's IRI, and sustaining ECP software and technologies critical to HCP operations and users.

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
<i>High Performance Data Facility, OPC</i> \$2,003	\$7,000	+\$4,997
Funding supports planning and preconceptual R&D for the HPDF, including site selection and preliminary design activities.	The Request will support pre-conceptual R&D and conceptual design for the HPDF project to support the site selection and Analysis of Alternatives processes in preparation for CD-1, and also potentially commencement of site preparation, contingent on achievement of CD-1 in FY 2024.	Increased funding will support the next step in design and planning of the HPDF project, advancing from pre-conceptual design to conceptual design, site selection, and alternative selection, and potentially, commencement of site preparation.
Leadership Computing Facilities \$430,000	\$466,607	+\$36,607
Funding supports operations at the LCF facilities at ANL and ORNL, including user support, power, space, system leases, early access systems and testbeds, and operations staff. Funding supports operations and allocation of exascale systems at OLCF and ALCF.	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 staff. The Request will support operations and allocation of exascale systems at OLCF and ALCF as well as planning for future upgrades, vendor partnerships, and DOE's IRI. The LCFs will deploy and maintain ECP software and technologies critical to HPC operations and users, and support ECP applications as they complete project milestones as ECP concludes.	Funding will support increased operating costs and system leases at both OLCF and ALCF to support operation of the exascale systems. Increase also supports planning for future upgrades, vendor partnerships, and DOE's IRI, as well as maintenance of ECP software and technologies critical to HPC operations and users.

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
<i>Leadership Computing Facility at ANL</i> \$175,000	\$219,000	+\$44,000
<p>Funding continues 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 deployed in calendar year 2022 and provides access for early science applications and the Exascale Computing Project. Competitive allocation of Aurora begins through ALCC for some exascale ready teams.</p>	<p>The Request will support start of operations and competitive allocation of the ALCF-3 exascale system, Aurora, which will deploy and maintain ECP software and technologies critical to HPC operations and users, and support ECP applications as they complete project milestones and the completion of the Exascale Computing Project. The Request will also support continuing operation and competitive allocation of the Polaris systems and other advanced computing testbeds. The Theta system will be decommissioned at the end of calendar year 2024. Increase supports significant increases in operating costs and lease payments for the Aurora exascale system as well as planning for future upgrades, vendor partnerships, and DOE's IRI.</p>	<p>Funding will support significantly increased operating costs and system lease payments for the Aurora exascale system, including power, maintenance, and space costs. Increase also supports planning for future upgrades, vendor partnerships, and DOE's IRI, as well as maintenance of ECP software and technologies critical to HPC operations and users.</p>
<i>Leadership Computing Facility at ORNL</i> \$255,000	\$247,607	-\$7,393
<p>Funding supports operations at the OLCF facility, including user support, power, space, system leases, maintenance, and staff. Funding also supports full operation and competitive allocation of the Frontier exascale system, Summit, and other testbeds.</p>	<p>The Request will support operations at the OLCF facility, including user support, power, space, maintenance, and staff. The Request will also support operation and competitive allocation of the Frontier exascale system and other advanced computing testbeds. OLCF will deploy and maintain ECP software and technologies critical to HPC operations and users, and support ECP applications as they complete project milestones as ECP concludes. Summit will be decommissioned at the end of calendar year 2024. Planning for OLCF-6 will begin, including vendor engagements. The Request also supports DOE's IRI.</p>	<p>Funding will support operating costs for the Frontier exascale system. Also, funding will support planning for future upgrades, vendor partnerships, and DOE's IRI, as well as maintenance of ECP software and technologies critical to HPC operations and users.</p>

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
High Performance Network Facilities and Testbeds	\$90,000	\$90,213 +\$213
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. Funding continues development of advanced network services at the start of operations of the recently completed 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.	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. Funding also supports planning and implementation of DOE's IRI.	Funding supports operations of ESnet and planning and implementation of DOE's IRI.

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, excluding facility operations.

Advanced Scientific Computing Research Exascale Computing Project and High Performance Data Facility

Description

SC and NNSA will complete 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 exascale systems at the LCFs, beginning in CY 2021 enabled the completion of ECI. As the project completes remaining KPPs, documentation and close out activities in FY 2024, ASCR will initiate site selection for a new HPDF. The HPDF will provide a crucial resource to SC programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources. The facility will be designed to dynamically configure computation, network resources and storage to access data at rest or in motion, supporting the use of well-curated datasets as well as near real-time analysis on streamed data directly from experiments or instruments.

24-SC-20, High Performance Data Facility

The FY 2024 Request includes \$1,000,000 in Total Estimated Cost (TEC) funding for the HPDF. The preliminary Total Project Cost (TPC) range for this project is \$300,000,000 to \$500,000,000. CD-0, Approve Mission Need, was approved on August 19, 2020. At that time, the scope of the project was broadly defined to include the potential for site preparation; construction or major upgrade of a data center facility; procurement of non-capital high performance computing, data storage, and local networking equipment; and non-recurring engineering activities with vendor partners to develop critical hardware and software components. Since CD-0, the scope of the project has evolved to consider a “Hub and Spoke” model of one Hub (the primary data center location) connected to several distributed Spokes located at key SC User Facilities and DOE national laboratories to address mission essential streaming data and edge applications as a critical enabler of DOE’s Integrated Research Infrastructure.

As early as 2013, a subcommittee of the Advanced Scientific Computing Advisory Committee (ASCAC) cited the need for a Data Facility in its transmittal report noting that “(1) a data-intensive storage and analysis facility with common interfaces and workflows will be necessary, and that (2) building on present ASCR facilities, at least in the near-term, will provide both early successes—such as NERSC’s work with Joint Genome Institute (JGI)—and considerable economies. In addition, there is often considerable synergy between analysis and visualization of large computational and observational data sets.”

With the resurgence of AI/ML and explosion of data volumes and velocities at many scientific user facilities, SC programs and their Scientific User Facilities have proposed accelerating discovery by developing new techniques to steer experiments and facilities; creating computing environments that integrate heterogeneous data for novel analyses; automating and streamlining interpretation of datasets; and making data Findable, Accessible, Interoperable, and Reusable (the FAIR principles of open data). These goals require new designs of computing and data infrastructure that provide researchers with reliable, simple, seamless performance and alleviate burdens from User Facility staff.

17-SC-20, SC Exascale Computing Project

The SC 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 what 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 DOE 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, 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 2024 Request includes \$14,000,000 for the SC-ECP. These funds will provide for delivery and documentation of remaining project milestones, documentation and open source publishing of the ECP software and technologies, transferring these technologies to the DOE computing facilities and to industry partners, and project close out activities including documenting application and software results and lessons learned.

**Advanced Scientific Computing Research
Exascale Computing Project and High Performance Data Facility**

Activities and Explanation of Changes

(dollars in thousands)

FY 2023 Enacted	FY 2024 Request	Explanation of Changes FY 2024 Request vs FY 2023 Enacted
Construction	\$77,000	\$15,000
		-\$62,000
17-SC-20, SC Exascale Computing Project	\$77,000	\$14,000
		-\$63,000
Funding supports project management and final execution of applications and software technology to meet the specified KPPs that demonstrate the development of an exascale ecosystem, which is the target of the project.	The Request will support project management to close out the project and final activities of applications and software technology to document results. FY 2024 is the last year of funding for the project. The project will focus on delivering remaining milestones and transferring ECP software and technologies to DOE computing facilities and industry partners.	FY 2023 was the final year of funding for the ECP applications and software teams. The funding will decrease to reflect the shift in focus from project execution to project close out.
24-SC-20, High Performance Data Facility	\$ —	\$1,000
		+\$1,000
No funding was appropriated in FY 2023 for this project.	The Request will support conceptual design for the HPDF project to support the site selection and Analysis of Alternatives processes in preparation for CD-1.	Funding will support the initial conceptual design activities.

**Advanced Scientific Computing Research
Capital Summary**

(dollars in thousands)

	Total	Prior Years	FY 2022 Enacted	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Capital Operating Expenses						
Capital Equipment	N/A	N/A	5,000	5,000	5,000	-
Total, Capital Operating Expenses	N/A	N/A	5,000	5,000	5,000	-

Capital Equipment

(dollars in thousands)

	Total	Prior Years	FY 2022 Enacted	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Capital Equipment						
Total, Non-MIE Capital Equipment	N/A	N/A	5,000	5,000	5,000	-
Total, Capital Equipment	N/A	N/A	5,000	5,000	5,000	-

Note:

- The Capital Equipment table includes MIEs located at a DOE facility with a Total Estimated Cost (TEC) > \$5M and MIEs not located at a DOE facility with a TEC >\$2M.

**Advanced Scientific Computing Research
Construction Projects Summary**

(dollars in thousands)

	Total	Prior Years	FY 2022 Enacted	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
24-SC-20, High Performance Data Facility						
Total Estimated Cost (TEC)	294,000	-	-	-	1,000	+1,000
Other Project Cost (OPC)	10,933	-	1,930	2,003	7,000	+4,997
Total Project Cost (TPC)	304,933	-	1,930	2,003	8,000	+5,997
17-SC-20, Exascale Computing Project (ECP)						
Total Estimated Cost (TEC)	695,376	517,376	115,000	63,000	-	-63,000
Other Project Cost (OPC)	630,830	588,830	14,000	14,000	14,000	-
Total Project Cost (TPC)	1,326,206	1,106,206	129,000	77,000	14,000	-63,000
Total, Construction						
Total Estimated Cost (TEC)	N/A	N/A	115,000	63,000	1,000	-62,000
Other Project Cost (OPC)	N/A	N/A	15,930	16,003	21,000	+4,997
Total Project Cost (TPC)	N/A	N/A	130,930	79,003	22,000	-57,003

**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.

(dollars in thousands)

	FY 2022 Enacted	FY 2022 Current	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Scientific User Facilities - Type A					
National Energy Research Scientific Computing Center	118,000	113,763	130,000	135,000	+5,000
Number of Users	8,500	9,110	9,200	9,200	–
Achieved Operating Hours	–	8,473	–	–	–
Planned Operating Hours	8,585	8,585	8,585	8,585	–
Unscheduled Down Time Hours	–	112	–	–	–
Argonne Leadership Computing Facility	160,000	154,233	175,000	219,000	+44,000
Number of Users	1,300	1,538	1,600	1,600	–
Achieved Operating Hours	–	6,987	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Unscheduled Down Time Hours	–	21	–	–	–
Oak Ridge Leadership Computing Facility	250,000	240,911	255,000	247,607	-7,393
Number of Users	1,500	1,674	1,700	1,700	–
Achieved Operating Hours	–	6,994	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Unscheduled Down Time Hours	–	14	–	–	–
Energy Sciences Network	90,000	86,715	90,000	90,213	+213
Number of Users	–	63	–	–	–
Achieved Operating Hours	–	8,760	–	–	–
Planned Operating Hours	8,760	8,760	8,760	8,760	–
Total, Facilities	618,000	595,622	650,000	691,820	+41,820
Number of Users	11,300	12,385	12,500	12,500	–

(dollars in thousands)

	FY 2022 Enacted	FY 2022 Current	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Achieved Operating Hours	-	31,214	-	-	-
Planned Operating Hours	31,361	31,361	31,361	31,361	-
Unscheduled Down Time Hours	-	147	-	-	-

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 2022 Enacted	FY 2023 Enacted	FY 2024 Request	FY 2024 Request vs FY 2023 Enacted
Number of Permanent Ph.Ds (FTEs)	819	825	813	-12
Number of Postdoctoral Associates (FTEs)	356	365	341	-24
Number of Graduate Students (FTEs)	523	535	595	+60
Number of Other Scientific Employment (FTEs)	217	220	182	-38
Total Scientific Employment (FTEs)	1,915	1,945	1,931	-14

Note:

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

**24-SC-20, High Performance Data Facility
Undesignated Laboratory
Project is for Design and Construction**

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

Summary

The FY 2024 Request for the Office of Science (SC) High Performance Data Facility (HPDF) project is \$1,000,000 of Total Estimated Cost (TEC) and \$7,000,000 of Other Project Costs (OPC). The preliminary Total Project Cost (TPC) range for this project is \$300,000,000 to \$500,000,000. The preliminary TPC estimate for this project is \$304,933,000.

The HPDF will provide a crucial resource to SC programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources. The facility will be designed to dynamically configure computation, network resources and storage to access data at rest or in motion, supporting the use of well-curated datasets as well as near real-time analysis on streamed data directly from experiments or instruments.

Significant Changes

This is a new Construction Project Data Sheet (CPDS) and this project is a new start in FY 2024. The most recent Department of Energy (DOE) Order 413.3B approved Critical Decision (CD) is CD-0, Approve Mission Need for a construction project with a conceptual scope and cost range, which was approved on August 19, 2020.

During FY 2023, the site for the HPDF will be selected and a Federal Project Director with the appropriate certification level will be assigned. The FY 2024 Request will support and conceptual design for the HPDF project to support the site selection and Analysis of Alternatives processes in preparation for CD-1, and also potentially commencement of site preparation, contingent on achievement of CD-1 in FY 2024.

Critical Milestone History

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	CD-2	Final Design Complete	CD-3	CD-4
FY 2024	8/19/20	3Q FY 2024	4Q FY 2024	4Q FY 2025	3Q FY 2025	4Q FY 2025	4Q FY 2030

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.

Project Cost History

(dollars in thousands)

Fiscal Year	TEC, Design	TEC, Construction	TEC, Total	OPC, Except D&D	OPC, Total	TPC
FY 2024	4,000	290,000	294,000	10,933	10,933	304,933

Note:

- This project has not received CD-2 approval; therefore, funding estimates are preliminary.

2. Project Scope and Justification

Scope

At CD-0 the scope of the project was broadly defined to include the potential for site preparation; construction or major upgrade of a data center facility; procurement of non-capital high performance computing (HPC), data storage, and local networking equipment; and non-recurring engineering activities with vendor partners to develop critical hardware and software components. Since CD-0, the scope of the project has evolved to consider a “Hub and Spoke” model of one Hub (the primary data center location) connected to several distributed Spokes located at key SC User Facilities and Department of Energy (DOE) national laboratories to address mission essential streaming data and edge applications as a critical enabler of DOE’s Integrated Research Infrastructure.

Justification

As early as 2013, a subcommittee of the Advanced Scientific Computing Advisory Committee (ASCAC) cited the need for a Data Facility in its transmittal report noting that “(1) a data-intensive storage and analysis facility with common interfaces and workflows will be necessary, and that (2) building on present Advanced Scientific Computing Research facilities, at least in the near-term, will provide both early successes—such as National Energy Research Scientific Computing Center’s work with Joint Genome Institute (JGI)—and considerable economies. In addition, there is often considerable synergy between analysis and visualization of large computational and observational data sets.”

With the growth of Artificial Intelligence and Machine Learning (AI/ML) and explosion of data volumes and velocities at many scientific user facilities, SC programs and their Scientific User Facilities have proposed accelerating discovery by developing new techniques to steer experiments and facilities; creating computing environments that integrate heterogeneous data for novel analyses; automating and streamlining interpretation of datasets; and making data Findable, Accessible, Interoperable, and Reusable (the FAIR principles of open data). These goals require new designs of computing and data infrastructure that provide researchers with reliable, simple, seamless performance and alleviate burdens from User Facility staff. Recent SC workshop reports and requirements reviews cite a number of challenges: Interaction with experiments in real time requires a service type that existing facilities do not provide such as the ability to guarantee a computing resource and quality of service during an experiment. AI/ML also requires the confluence of large well-curated datasets and the compute resources to perform net training activities. Currently, most analyses of experimental and simulation data are done post hoc, after the experiment or simulation has run. Controlling either extreme-scale simulation or experimental facilities with AI requires low-latency analysis and inference using high-volume, high-velocity data sets in real time. Traditional HPC systems are designed to efficiently execute large-scale simulations and focused on minimizing users’ wait-times in batch queues. The SC Integrated Research Infrastructure Architecture Blueprint Activity, a convening of over 160 DOE laboratory subject matter experts, identified the need for new high performance data infrastructure to advance these goals as part of a DOE’s Integrated Research Infrastructure vision.

The proposed HPDF will serve as a foundational element in enabling the DOE Integrated Research Infrastructure; will provide crucial resources to Office of Science programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources; will partner and operate in concert with other ASCR Facilities and potentially other DOE laboratory computing resource providers to provide a high availability high performance computing ecosystem for a wide variety of applications; will serve as a “Hub” enabling “Spoke” sites to deploy and orchestrate distributed infrastructure to enable high priority DOE mission applications.

The project is being conducted in accordance with the project management requirements in DOE Order 413.3B, Program and Project Management for the Acquisition of Capital Assets.

Key Performance Parameters (KPPs)

The KPPs will be determined after the site selection process and the analysis of alternatives has completed, expected around 3Q FY 2024, and upon establishment of the conceptual design prior to achievement of CD-1.

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Design (TEC)			
FY 2024	1,000	1,000	500
Outyears	3,000	3,000	3,500
Total, Design (TEC)	4,000	4,000	4,000
Construction (TEC)			
Outyears	290,000	290,000	290,000
Total, Construction (TEC)	290,000	290,000	290,000
Total Estimated Cost (TEC)			
FY 2024	1,000	1,000	500
Outyears	293,000	293,000	293,500
Total, TEC	294,000	294,000	294,000

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Other Project Cost (OPC)			
FY 2022	1,930	1,930	–
FY 2023	2,003	2,003	–
FY 2024	7,000	7,000	7,933
Outyears	–	–	3,000
Total, OPC	10,933	10,933	10,933

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Project Cost (TPC)			
FY 2022	1,930	1,930	–
FY 2023	2,003	2,003	–
FY 2024	8,000	8,000	8,433
Outyears	293,000	293,000	296,500
Total, TPC	304,933	304,933	304,933

4. Details of Project Cost Estimate

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Design	2,600	N/A	N/A
Design - Contingency	1,400	N/A	N/A
Total, Design (TEC)	4,000	N/A	N/A
Construction	188,500	N/A	N/A
Construction - Contingency	101,500	N/A	N/A
Total, Construction (TEC)	290,000	N/A	N/A
Total, TEC	294,000	N/A	N/A
<i>Contingency, TEC</i>	<i>102,900</i>	<i>N/A</i>	<i>N/A</i>
Other Project Cost (OPC)			
OPC, Except D&D	7,106	N/A	N/A
OPC - Contingency	3,827	N/A	N/A
Total, Except D&D (OPC)	10,933	N/A	N/A
Total, OPC	10,933	N/A	N/A
<i>Contingency, OPC</i>	<i>3,827</i>	<i>N/A</i>	<i>N/A</i>
Total, TPC	304,933	N/A	N/A
Total, Contingency (TEC+OPC)	106,727	N/A	N/A

5. Schedule of Appropriations Requests

(dollars in thousands)

Fiscal Year	Type	Prior Years	FY 2022	FY 2023	FY 2024	Outyears	Total
FY 2024	TEC	—	—	—	1,000	293,000	294,000
	OPC	—	1,930	2,003	7,000	—	10,933
	TPC	—	1,930	2,003	8,000	293,000	304,933

6. Related Operations and Maintenance Funding Requirements

Start of Operation or Beneficial Occupancy	4Q FY 2030
Expected Useful Life	TBD
Expected Future Start of D&D of this capital asset	4Q FY 2036

Related Funding Requirements
(dollars in thousands)

	Annual Costs		Life Cycle Costs	
	Previous Total Estimate	Current Total Estimate	Previous Total Estimate	Current Total Estimate
Operations	N/A	TBD	N/A	TBD
Utilities	N/A	TBD	N/A	TBD
Maintenance and Repair	N/A	TBD	N/A	TBD
Total, Operations and Maintenance	N/A	TBD	N/A	TBD

Notes:

- *The project is likely to comprise both capital assets (refurbishment or build of data center space) and non-capital assets (IT components that comprise the computational and data management infrastructure). The expected useful life of the former is potentially 10–20 years, while the latter is 5–7 years.*
- *Life-Cycle costs will be performed as part of CD-1.*

7. D&D Information

The scope and nature of D&D activities will be determined at CD-1.

	Square Feet
New area being constructed by this project at [Lab]	TBD
Area of D&D in this project at [Lab]	TBD
Area at [Lab] to be transferred, sold, and/or D&D outside the project, including area previously “banked”	TBD
Area of D&D in this project at other sites	TBD
Area at other sites to be transferred, sold, and/or D&D outside the project, including area previously “banked”	TBD
Total area eliminated	TBD

8. Acquisition Approach

Once the site is selected, SC will work with the selected laboratory to determine if a building to house the facility will need to be constructed. All computing and storage resources, as well as, non-recurring engineering activities will be procured through open solicitations.

**17-SC-20, SC Exascale Computing Project
Oak Ridge National Laboratory, ORNL
Project is for Design and Construction**

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

Summary

The FY 2024 Request for the SC Exascale Computing Project (SC-ECP) is \$14,000,000 of Other Project Costs (OPC). 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 (10¹⁸ 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.

Significant Changes

This project was initiated in FY 2017. The FY 2024 Request supports investments in the ECP technical focus areas—application development, software technology and hardware and integration—to support the final close out of the activities which includes developing technical writeups, hardening software and collection of lessons learned. The funding decrease reflects the achievement, in FY 2023, of the threshold Key Performance Parameters (KPPs) as well as a subset of the objective KPPs.

Critical Milestone History

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	CD-2	Final Design Complete	CD-4
FY 2024	7/28/16	3/22/16	1/3/17	2/25/20	6/6/19	4Q FY 2024

Fiscal Year	Performance Baseline Validation	CD-3A	CD-3B
FY 2024	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 2023	—	700,843	700,843	625,363	625,363	1,326,206
FY 2024	—	695,376	695,376	630,830	630,830	1,326,206

2. Project Scope and Justification

Scope

Four well-known challenges^a are key to requirements and Mission Need of the SC-ECP. These challenges are:

- Parallelism: Systems must exploit the extreme levels of parallelism that will be incorporated in an exascale-capable computer;
- Resilience: Systems must be resilient to permanent and transient faults;
- Energy Consumption: System power requirements must be no greater than 20-30 MW; and
- Memory and Storage Challenge: Memory and storage architectures must be able to access and store information at anticipated computational rates.

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.

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 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,

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

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

Key Performance Parameters (KPPs)

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

Performance Measure	Threshold	Objective
Exascale performance improvements for mission-critical challenge problems	50 percent of selected applications achieve Figure of Merit improvement greater than or equal to 50x	100 percent of selected applications achieve their KPP-1 stretch goal
Broaden exascale science and mission capability	50 percent of the selected applications can execute their challenge problem ^a	100 percent of selected applications can execute their challenge problem stretch goal
Productive and sustainable software ecosystem	50 percent of the weighed impact goals are met	100 percent of the weighted impact goals are met
Enrich the HPC Hardware Ecosystem	Vendors meet 80 percent of all the PathForward milestones	Vendors meet 100 percent of all the PathForward milestones

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

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Construction (TEC)			
Prior Years	517,376	517,376	257,725
FY 2022	115,000	115,000	163,147
FY 2023	63,000	63,000	221,745
FY 2024	–	–	47,292
Outyears	–	–	5,467
Total, Construction (TEC)	695,376	695,376	695,376
Total Estimated Cost (TEC)			
Prior Years	517,376	517,376	257,725
FY 2022	115,000	115,000	163,147
FY 2023	63,000	63,000	221,745
FY 2024	–	–	47,292
Outyears	–	–	5,467
Total, TEC	695,376	695,376	695,376

Note:

- The project approved a project change request to extend technical campaigns one quarter into FY 2024 to allow full access to the Aurora computer which was delayed due to supply chain issues currently prevalent since COVID. This will not impact the CD-4 baseline date.

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Other Project Cost (OPC)			
Prior Years	588,830	588,830	583,920
FY 2022	14,000	14,000	9,171
FY 2023	14,000	14,000	18,753
FY 2024	14,000	14,000	13,983
Outyears	–	–	5,003
Total, OPC	630,830	630,830	630,830

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Project Cost (TPC)			
Prior Years	1,106,206	1,106,206	841,645
FY 2022	129,000	129,000	172,318
FY 2023	77,000	77,000	240,498
FY 2024	14,000	14,000	61,275
Outyears	–	–	10,470
Total, TPC	1,326,206	1,326,206	1,326,206

4. Details of Project Cost Estimate

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

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total Estimated Cost (TEC)			
Application Development (TEC)	347,349	347,349	346,360
Production Ready Software	228,356	228,356	217,290
Hardware Partnership	125,138	125,138	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	89,689	89,688
Application Development (OPC)	221,050	221,050	221,050
Software Research	118,517	118,517	118,517
Hardware Research	196,107	196,107	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 2022	FY 2023	FY 2024	Outyears	Total
FY 2023	TEC	522,843	115,000	63,000	—	—	700,843
	OPC	583,363	14,000	14,000	—	14,000	625,363
	TPC	1,106,206	129,000	77,000	—	14,000	1,326,206
FY 2024	TEC	517,376	115,000	63,000	—	—	695,376
	OPC	588,830	14,000	14,000	14,000	—	630,830
	TPC	1,106,206	129,000	77,000	14,000	—	1,326,206

6. Related Operations and Maintenance Funding Requirements

System procurement activities for the exascale-capable computers are not part of the SC-ECP. The exascale-capable computers will become part of existing facilities and operations and maintenance funds, and will be included in the ASCR facilities’ operations or research program’s budget. A Baseline Change Proposal (BCP) was executed in March 2018 to reflect this change.

Start of Operation or Beneficial Occupancy	4Q FY 2024
Expected Useful Life	7 years
Expected Future Start of D&D of this capital asset	4Q FY 2031

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.