Advanced Scientific Computing Research

Overview

The Advanced Scientific Computing Research (ASCR) program's mission is to advance applied mathematics and computer science; deliver the most sophisticated computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and software tools for science and engineering in partnership with the research community, including U.S. industry. ASCR supports state-of-the-art capabilities that enable scientific discovery through computation. 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 other 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, including the unique expertise needed at the forefront of this strategic technology.

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.[a] 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 HPC and trustworthy 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 also 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 science will come from employing data-driven methods in AI at extreme scales 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. ASCR is in a pivotal position to leverage the exascale ecosystem and decades of basic research investments as well as industry partnerships to drive responsible development of AI technologies and AI enabled science in the national interest.

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.

Continued progress in the scientific utilization of microelectronics, especially the energy utilization of these devices for HPC and AI, underpins all ASCR's efforts. ASCR's strategy is to focus on technologies that build on expertise and core investments across SC, continuing mutually beneficial engagements with industry, the applied technology offices, other agencies, and the scientific community through connections made in the Exascale Computing Project (ECP); investing in small-scale testbeds; and increasing core research investments in Applied Mathematics and Computer Science.

---

[a] https://nap.nationalacademies.org/catalog/21886/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-2017-2020

ASCR's proposed activities will deliver on the promise of the exascale and AI-enabled science era 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 2025 Request

The FY 2025 Request for $1,152.7 million for ASCR is an increase of $84.7 million over the FY 2023 Enacted, and is well-aligned with Administration and Department priorities to advance responsible AI technology, critical and emerging technologies such as QIS and microelectronics. In addition, these investments continue to address the challenges of climate change to reach netzero through the Earthshots awards and contribute to better health outcomes in pandemic readiness through continued support for Cancer Moonshot. It also provides support to reduce barriers and inequities through workforce investments, facilitate adoption of accelerated HPC, and usher in the responsible AI and exascale science era to bolster industrial innovation.

<u>Research</u>

- The Request prioritizes delivering on the promise of the exascale and AI enabled science era including critical basic research investments in applied mathematics and computer science to merge the power of AI with exascale computing, develop tools that facilitate building foundation models useful for basic and applied science, and partnerships that build and use foundation models supporting new applications in science and energy, national security, and increased community preparedness. The Request also emphasizes applied mathematics, computer science, networking, hardware, and microelectronics research to advance and leverage energy-efficient advanced computing including quantum. Investments in support of the Energy Earthshot Research Centers provide the research underpinning the DOE's stretch goals. Increased or shifted efforts in both research and at the facilities will advance implementation of DOE's Integrated Research Infrastructure (IRI) to integrate DOE's unique data, user facilities, and computing resources. SC Microelectronics Science Research Centers will comprise a network of multiple team awards, with individual awards focused on a dimension related to a common research topic for each Center. The multidisciplinary teams will include researchers from universities, national laboratories, and industry. Materials, chemistries, devices, systems, architectures, algorithms, and software will be developed in a closely integrated fashion in a co-design innovation ecosystem. Strategic partnerships, both within DOE and at the interagency level, expand the impact of the ECP, AI, and accelerate scientific discovery through advanced computing (SciDAC), including to support national emergency preparedness and improved health (BRaVE). Underpinning all investments are increases in efforts to grow the necessary workforce through Reaching a New Energy Sciences Workforce (RENEW), FAIR, Computational Sciences Graduate Fellowship (CSGF), and Established Program to Stimulate Competitive Research (EPSCoR).

- 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. This support enables the recompetition/renewal of the NQISRCs and expansion of 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 provide national leadership in quantum in coordination with relevant agencies.

<u>Facility Operations</u>

- FY 2025 Request supports increases for operations and competitive allocation of the Nation's exascale computing systems: Frontier at the Oak Ridge Leadership Computing Facility (OLCF), deployed in calendar year 2021; and Aurora at Argonne Leadership Computing Facility (ALCF), deployed in calendar year 2023. Increased 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 maintain ECP software and technologies critical to HPC operations and users. In addition, increased funding supports: planning for NERSC-10, OLCF-6 and ALCF-4—including site preparations, long lead procurements, and vendor R&D partnerships—to address rising demand for computing and U.S. competitiveness in HPC and computational science; operation and expansion of IRI; and construction of a new High Performance Data Facility (HPDF), to strengthen and leverage SC's unparalleled research capabilities.

- The Request supports new IRI efforts to prioritize development of state-of-the-art real-time experimental/observational workflows and drive innovation in system architectures and services to inform planning for upgrades at the Facilities.

Projects

- The ECP project was successfully concluded in FY 2024 and no funding is included in the FY 2025 Request.
- The FY 2025 Request supports HPDF final design planning to establish cost, scope, and schedule in preparation for CD-2/3.

Advanced Scientific Computing Research
Funding

(dollars in thousands)

	FY 2023 Enacted	FY 2024 Annualized CR	FY 2025 Request	FY 2025 Request vs FY 2023 Enacted
Advanced Scientific Computing Research				
Applied Mathematics Research	61,035	60,438	77,565	+16,530
Computer Sciences Research	60,667	86,267	86,736	+26,069
Computational Partnerships	95,875	77,405	93,449	-2,426
Advanced Computing Research	108,920	128,598	148,197	+39,277
Energy Earthshot Research Centers	12,500	15,000	12,500	–
Total, Mathematical, Computational, and Computer Sciences Research	**338,997**	**367,708**	**418,447**	**+79,450**
High Performance Production Computing	132,003	136,000	146,500	+14,497
Leadership Computing Facilities	430,000	441,000	475,195	+45,195
High Performance Network Facilities and Testbeds	90,000	88,400	93,540	+3,540
Integrated Research Infrastructure	–	–	3,000	+3,000
Total, High Performance Computing and Network Facilities	**652,003**	**665,400**	**718,235**	**+66,232**
17-SC-20, SC Exascale Computing Project	**77,000**	**–**	**–**	**-77,000**
Subtotal, Advanced Scientific Computing Research	**1,068,000**	**1,033,108**	**1,136,682**	**+68,682**
Construction				
24-SC-20 High Performance Data Facility	–	–	16,000	+16,000
Subtotal, Construction	**–**	**–**	**16,000**	**+16,000**
Total, Advanced Scientific Computing Research	**1,068,000**	**1,033,108**	**1,152,682**	**+84,682**

SBIR/STTR funding:

- FY 2023 Enacted: SBIR $10,112,000 and STTR $1,422,000
- FY 2024 Annualized CR: SBIR $10,775,000 and STTR $1,515,000
- FY 2025 Request: SBIR $12,046,000 and STTR $1,694,000

Advanced Scientific Computing Research
Explanation of Major Changes

	(dollars in thousands) FY 2025 Request vs FY 2023 Enacted
Mathematical, Computational, and Computer Sciences Research	**+$79,450**
The Request significantly increases investments in trustworthy and responsible AI to develop tools that facilitate building foundation models useful for basic and applied science, including expanded partnerships with industry, academia, other agencies, and international allies. DOE will utilize its computing capabilities, AI testbeds, and research efforts to increase participation and retention of emerging research institutions and next generation researchers through EPSCoR, CSGF, FAIR and RENEW. Computer Science and Applied Mathematics activities will continue foundational and long-term basic research efforts that: explore and prepare for emerging technologies; develop new scalable energy efficient algorithms and software; address the challenges of data intensive science and emerging computing technologies, such as quantum information science; support the development of safe, secure, and trustworthy AI technologies and associated Privacy Enhancing Technologies; support a network of multiple team awards for the Microelectronics Science Research Centers, with individual awards focused on a dimension related to a common research topic for each Center; and will advance the IRI. Computational Partnerships supports partnerships across DOE and with other agencies to expand the impact of ECP and advance national goals with funding reduced to better align IRI efforts. The Advanced Computing Research activity will support the recompetition/renewal of the NQISRCs and quantum computing and networking testbeds, in close coordination with the other SC programs.	
High Performance Computing and Network Facilities	**+$66,232**
The increase prioritizes OLCF and ALCF to 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, including AI. In addition, funding supports the operation of the 125 petaflop NERSC-9 Perlmutter system. The increase also supports operations of all facilities—including power and cooling, equipment, staffing, testbeds, lease payments, user programs, outreach, and continued implementation of SC's IRI— as well as planning, site preparations and project efforts for NERSC-10 and LCF upgrades.	
Exascale Computing	**-$77,000**
The ECP was successfully completed in FY 2024.	
Construction	**+$16,000**
The FY 2025 Request supports the HPDF for the final design, and to establish the project cost, scope, and schedule.	
Total, Advanced Scientific Computing Research	**+$84,682**

Basic and Applied R&D Coordination

Coordination across disciplines and programs is a cornerstone of the ASCR program. Partnerships within SC and National Nuclear Security Administration (NNSA) continue in advanced computing and applications. ASCR also has partnerships in QIS and AI within SC and is collaborations across DOE and with other agencies to expand the AI-enabled Exascale science era. Through the Networking and Information Technology R&D Subcommittee of the National Science and Technology Council (NSTC) Committee on Technology, ASCR coordinates with programs across the Federal Government. Future Advanced Computing technologies, Scientific Data, Large Scale Networking, High End Computing, AI, and QIS are coordinated with other agencies through the NSTC. In FY 2025, cross-agency interactions and collaborations continue in coordination with the Office of Science and Technology Policy.

Program Accomplishments

Delivering a Capable Exascale Computing Ecosystem for the Nation

The ECP met all of the project's key performance parameters in 2023 and successfully completed, documented and closed out the project in FY 2024, under budget and ahead of schedule. The project enabled the launch of exascale systems that use less than 20MW and created a modern, interoperable, and portable software ecosystem that addresses the needs of simulations, big data, and AI with numerous awards including 13 R&D 100 awards, 7 Gordon Bell finalists and prizes, and dozens of industry recognitions. For example, ECP's Extreme Scale Scientific Software Stack (E4S) provides easy access to over 100 HPC, AI, and data analytics packages and tools—all of which are ready for deployment for a variety of accelerated heterogenous HPC architectures, including systems with Graphical Processing Units (GPUs) from multiple vendors. This centralized, standardized, build-and-deploy framework addresses software issues and saves both time and money. It also, makes ECP software available to researchers across the Nation and lowers barriers to expanding user access to HPC and ECP software. These investments de-risked the jump to accelerated HPC for a wide range of use cases in science and engineering. As a result, scores of companies, universities, and government labs from across the Nation have stood up compatible hardware and installed ECP software and applications - with many more leverging ECP technologies through cloud service providers.[b]

Harnessing Exascale Toward Solving the Global Climate Crisis

The Gordon Bell Prize is a prestigious award given each year to recognize outstanding achievement in HPC. Starting in 2023, an additional Gordon Bell Prize for Climate Modeling will be awarded every year, for ten years, to recognize innovative parallel computing contributions toward solving the global climate crisis. The ECP's Energy Exascale Earth System Model (E3SM), a partnership with Biological and Environmental Research, was awarded the inaugural prize for work on an efficient and performance-portable implementation of a Simple Cloud Resolving E3SM Atmosphere Model (SCREAM). E3SM submitted performance results on Frontier, the first exascale computer, and benchmarked the model configuration used for scientific research. The paper outlines results which represent several firsts for a global cloud-resolving model (GCRM): first GCRM to run on an Exascale supercomputer, first GCRM to run at scale on both NVIDIA and AMD GPU systems, and first nonhydrostatic GCRM to exceed one simulated-year-per-day (SYPD) of full model throughput. SCREAM is a monumental advance that avoids the uncertainties and biases often associated with lower resolution models. For example, it captures the structure of important weather events, such as cyclones, atmospheric rivers, and cold air outbreaks, which are poorly captured by typical global cloud models. This marks the beginning of a revolution in DOE earth system modeling where unprecedentedly detailed and realistic simulations lead the way to more accurate predictions.

Industry Partners Leverage ECP to Launch Exascale Innovation Era

One of the hallmarks of ECP has been collaboration with industry partners who have helped promote exascale computing and the successful integration of software technologies, applications, and hardware to deliver a fully capable exascale ecosystem. For industry partners, exascale capabilities open pathways to explore promising new technologies and products and reduce time to commercialization. For example, GE Research is using Frontier for virtual testing of innovative new open fan jet engine designs to achieve greater fuel efficiency and advance their goal to reduce CO2 emissions by 20 percent over today's most efficient engines. These simulations are computationally intensive, requiring significant resources to resolve the necessary length and timescales. With Frontier, researchers can simulate the movement of air and complex flow characteristics at full flight-scale conditions for the first time and reveal previously unresolvable details of operation. As a result, researchers gain insight into realistic performance from design options years ahead of building the physical system.

---

[b] https://www.exascaleproject.org/industry-and-agency-council/

Research investments can then be focused on the most promising designs, accelerating time to production and resulting in substantial cost-savings. This work also lays the foundation for design improvements crucial to advancing jet engine design toward alternative fuels like hydrogen and sustainable aviation fuels as well as new hybrid-electric technologies.

Interagency Partners Leverage Exascale to Improve Cancer Outcomes
ECP included a unique partnership with the National Cancer Institute as part of the Cancer Moonshot. A focus on this effort was to develop CANDLE (CANcer Distributed Learning Environment), an AI-based computer code that brings together DOE capabilities and cancer research to accelerate discovery of new cancer therapies and improve outcomes. This platform aims to address three challenges: understanding protein behaviors in tumor cells; understanding the relationship between tumors and drugs; and analyzing population-level biomedical records to extract new patterns and information. To address the third challenge, the Modeling Outcomes Using Surveillance Data and Scalable AI for Cancer (MOSSAIC) project, which applies natural language processing and deep learning algorithms, with uncertainty quantification and privacy preserving capabilities, to population-based cancer data from NCI's Surveillance, Epidemiology, and End Results (SEER) program. MOSSAIC efficiently reduced manual efforts while improving accuracy in the NCI SEER program, cutting costs for the SEER registries, and providing a trustworthy source of population-level cancer surveillance data, at closer to real-time, for use by the cancer research community. MOSSAIC is also leveraging OLCF's Frontier exascale system to develop privacy preserving federated learning tools that have the potential to substantially increase secure AI model development and sharing across the cancer community. The CANDLE framework was recognized with a 2023 R&D 100 Award and has been used to research potential treatments for SARS-CoV-2 and to assess the performance of DOE's exascale computing systems.

Accelerating Neuromorphic Computing
To address the current limitations and inefficiencies that inhibit large-scale neuromorphic computing, researchers at ORNL created SuperNeuro, a Python-based open software that provides AI practitioners with brain-like simulators that are fast and scalable on central and graphics processing platforms. Using matrix-based and agent-based modeling approaches, SuperNeuro allows for different workloads and provides the option of simulating the user's own spiking mechanisms in a human-interpretable manner. Compared with existing simulation platforms, SuperNeuro can be up to 300 times faster for small sparse networks and up to 3.4 times faster on large sparse and dense network, leveraging GPU computing to provide superior performance for neuroscience, spiking neural networks, or SNNs, and general-purpose computing workloads. SuperNeuro is also more adaptable and provides easy integration with other tools for SNN optimization that opens the possibilities for codesign of neuromorphic circuits. SuperNeuro was recognized with a 2023 R&D 100 award.

Proving Entanglement is Key to Quantum Advantage
For decades scientists have had theoretical evidence that quantum computers can solve problems that are too difficult for even today's exascale systems but the source of this advantage has been uncertain. Entanglement, a fundamental property of quantum systems, has been a key suspect but is a complex phenomenon that is difficult to pin down in today's noisy quantum systems. In work published in Physical Review Letters, a team of ASCR researchers have identified a computational problem that demonstrates a provable quantum speedup over any classical computation in which control of the quantum entanglement is directly responsible for causing the speedup. This research is a crucial step in realizing the potential of quantum computing by identifying the path to practical applications on real world system.

Democratizing Quantum Control
The quantum systems that hold great promise for scientific computing and networking are delicate, requiring sophisticated and often costly control electronics to achieve state-of the-art performance. This has erected a barrier to entry to newcomers to the field and limited the ability of even the best academic researchers to advance the frontiers of quantum information science. The Quantum Instrumentation Control Kit (QICK) board aims to change that. Initially developed at FNAL and brought to maturity with support from the ORNL-led Quantum Science Center, one of the five National Quantum Information Science Research Centers authorized by the 2018 National Quantum Initiative Act, the QICK board combines a commercial off-the-shelf programmable logic board with open-source software that can be adapted to any type of quantum experiment. Dozens of research groups in academia, National Laboratories, and industry are now using QICK board to advance their own research ranging from development of novel silicon-based qubits to qubit-based sensors for dark matter detection.

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 climate science, and energy-efficient microelectronics. 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 and help realize the promise of the exascale and AI-enabled science era. Research efforts 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. In part through continued funding for the EPSCoR, RENEW, and FAIR initiatives, ASCR will build stronger programs with underserved communities and emerging research institutions (ERIs) as well as HBCUs and MSIs, including investing in a more diverse and inclusive workforce.

Applied Mathematics Research

The FY 2025 Request for 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, methods that facilitate building foundational models for trustworthy and privacy preserving AI/ML, and efficient data analysis underpin all of DOE's computational and data-intensive science efforts. More broadly, the Request supports 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. Industry often uses software developed with Applied Mathematics investments and integrate it with their own software.

Computer Science Research

The FY 2025 Request for the Computer Science 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 incorporation of AI techniques and real-time exploration and understanding of extreme scale and complex data from both simulations and experiments. Additionally, Computer Science efforts play a key role in understanding gaps and future opportunities for the design of future computing systems that maintains U.S. leadership in high-performance, data-intensive, and AI computing. To support these goals, this activity includes support for foundational research in data analysis and visualization, data management and storage, distributed systems and resource management, programming models and tools enabling high performance and portability, program verification and testing, operating and runtime systems, advanced networking, hardware/software co-design, and energy-aware computer-science fundamentals. Hardware and software vendors often use software developed with ASCR Computer Science investments and integrate it with their own software. In addition, partnerships between mathematicians and computer scientists, jointly supported by this activity and Applied Mathematics, develops energy efficient algorithms and methods that scale from intelligent sensors to HPC to advance the Department's energy goals.

Computational Partnerships

The FY 2025 Request for 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, including AI/ML, 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. These efforts include partnerships with industry, academia, other agencies, and international allies to utilize DOE's computing capabilities and AI testbeds to build foundation models that support new applications in science and energy.

This FY 2025 Request for this activity also supports the FAIR initiative, which provide focused investment on enhancing research on clean energy, climate, and related topics, including attention to emerging research institutions (ERIs), HBCUs and MSIs, underserved regions and communities, 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.

Advanced Computing Research

This FY 2025 Request for the Advanced Computing Research activity supports efforts focused on development of emerging computing technologies such as QIS and neuromorphic computing 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 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.

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.

RENEW investments leverage SC's unique national laboratories, user facilities, and other research infrastructures to provide undergraduate and graduate training opportunities for emerging academic institutions and underserved Historically Black College and University (HBCU) or Minority Serving Institution (MSI) communities in the SC research portfolio.

SC Microelectronics Science Research Centers will comprise a network of multiple team awards, with individual awards focused on a dimension related to a common research topic for each Center. The multidisciplinary teams will include researchers from universities, national laboratories, and industry. Materials, chemistries, devices, systems, architectures, algorithms, and software will be developed in a closely integrated fashion in a co-design innovation ecosystem. The Request will support 2 to 3 ASCR research awards that would contribute to one of up to 4 cross-SC Research Centers.

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 FY 2025 Request support's ASCR's continued partnership with SC's Basic Energy Sciences (BES) and Biological and Environmental Research (BER) programs in the EERCs, a new modality of research launched in FY 2023, as well as directed fundamental research 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 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Mathematical, Computational, and Computer Sciences Research	**$338,997**	**$418,447**	**+$79,450**
Applied Mathematics Research	$61,035	$77,565	+$16,530
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 innovative research efforts in algorithms, libraries and methods that underpin high-end scientific simulations, trustworthy and privacy enhancing scientific AI/ML techniques including methods that facilitate building foundation models useful for basic and applied science, 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.	Funding will increase support for basic research that addresses foundational applied math challenges critical to enable science at the exascale era, and the development and evaluation of trustworthy and privacy-preserving AI methods integrating foundation models with task-specific data-driven capabilities.

(dollars in thousands)

FY 2023 Enacted	FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Computer Science Research $60,667	$86,736	+$26,069
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.	The Request will continue support for innovative 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 to facilitate building foundation models useful for basic and applied science and to enable an integrated research 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. In addition, funding will support partnerships between mathematicians and computer scientists to develop energy efficient scalable algorithms and methods.	Funding will increase support for basic research that addresses foundational computer science challenges critical to enable science at the exascale era and the development and evaluation of federated, trustworthy, privacy enhancing, AI methods integrating foundation models with task-specific data-driven capabilities.

(dollars in thousands)

FY 2023 Enacted		FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Computational Partnerships	$95,875	$93,449	-$2,426
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.		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, AI/ML, 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 initiative.	Decrease reflects minor adjustments in the Advanced Computing portfolio to better align with ongoing Integrated Research Infrastructure initiative.

(dollars in thousands)

FY 2023 Enacted		FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Advanced Computing Research	$108,920	$148,197	+$39,277
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.		The Request will continue to support quantum computing testbed efforts, and regional quantum internet testbeds. The Request allows REP to increase strategic investments in emerging technologies including AI-focused hardware, and continued support for hardening of critical software developed under ECP to enable science at the exascale era. Small investments in cybersecurity will continue. The Request will increase support for the CSGF fellowship, in partnership with NNSA, to support increased tuition costs and stipends, and to increase the number of fellows focused on AI, 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 the exascale and AI-enabled computational science era and beyond Moore's Law capabilities such as QIS. The Request will also continue support for the RENEW initiative to provide undergraduate and graduate training opportunities for students and academic institutions not currently well represented in the U.S. S&T ecosystem to expand the pipeline for ASCR research and facilities workforce needs. The NQISRCs will be recompeted, as authorized in the National Quantum Initiative Act. The Request will support 2 to 3 research awards that would contribute to one of up to 4 cross-SC Microelectronics Science Research Centers.	The Request will support increases for AI-focused hardware investments, Microelectronics Science Research Centers, and CSGF.NQISRCs will be recompeted.

(dollars in thousands)

FY 2023 Enacted		FY 2025 Request	Explanation of Changes FY 2025 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 continues to support the EERCs established jointly between Office of Science programs (BES, ASCR, and BER) with strong coordination the DOE Applied Technology Offices. EERC efforts will continue to inform foundational research investments in applied mathematics and computer science that address the longer-term challenges of the Energy Earthshots.	No change.

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 High Performance Production Computing (HPC) and Network Facilities subprogram supports the construction and operations of forefront research computing, networking, and data user facilities to meet critical mission needs. The HPC activity supports the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL), which provides HPC resources and large-scale storage to a broad range of SC researchers, and the High Performance Data Facility (HPDF) that will provide a managed computational and data resource to attack fundamental problems in science and engineering. The Leadership Computing activity supports the two Leadership Computing Facilities (LCFs) at Oak Ridge National Laboratory (ORNL) and Argonne National Laboratory (ANL), which provide diverse leading-edge HPC capabilities to the U.S. research and industrial communities. The High Performance Network Facilities and Testbeds activity supports the high performance network user facility, ESnet, which connects all DOE national laboratories and many sites to global research networks and delivers highly reliable data transport capabilities optimized for the requirements of large-scale science. Within the subprogram, facility operations include investments in upgrade projects. The core strength of the facilities is the dedicated staff who work to maximize user productivity and science impact, operate and maintain world-leading research computing, networking, and data infrastructure, while simultaneously executing major upgrade projects.

The HPC and Network Facilities subprogram investments are informed through formal collection of strategic user requirements for research computing and data management from stakeholders across SC and DOE, including the other SC research programs, SC scientific user facilities, DOE national laboratories, and other stakeholders. ASCR continues to observe an accelerating pace of innovation in computing technology through and beyond the exascale era.

Allocation of HPC resources to users follows the merit review public-access model used by all 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.

In FY 2025, the facilities will continue implementation of 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. At the dawn of the exascale science era, many researchers and collaborations strive to meld data, simulation, and AI tools in novel ways, some with strict operational demands. Agency and program leaders feel the urgency to bring the best-integrated science approaches to bear on our greatest challenges. Implementing the IRI vision requires the creation of an integrated research ecosystem that empowers researchers to rapidly accelerate time to insight.

<u>High Performance Production Computing</u>

The FY 2025 Request for 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 researchers conducting over 1,000 projects use NERSC annually to perform scientific research across a wide range of disciplines. 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 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.

In FY 2025, NERSC will operate the 125 pf HPE/AMD/NVIDIA NERSC-9 system (Perlmutter), an AI-enabled GPU-CPU system, 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. In addition, the diversity of data- and compute-intensive research workflows is expanding rapidly. As demand for HPC resources grows and diversifies, ASCR foresees the strategic need for operational resilience and software portability across its HPC resources. The FY 2025 Request also 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 and use cases. NERSC operations funding also contributes to expanding 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.

In FY 2024, the HPC activity will continue planning for the HPDF project, which will serve as a foundation for the IRI. HPDF will provide a managed computational and data 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 and 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, scaling tests, early science applications, and tool and library developers. Their efforts are critical enablers of partnerships to broaden the benefits of exascale computing for the Nation. Industry use of the LCFs, often prompt companies to expand their own HPC resources.

The FY 2025 Request for this activity supports operation and competitive allocation of the OLCF at ORNL, including the Nation's first exascale computing system, an HPE-Cray/AMD exascale system (Frontier), deployed in 2021. Funding also supports decommissioning of the 200 PF IBM/NVIDIA OLCF-4 system (Summit), expansion of the Quantum Computing User Program, IRI efforts, AI and advanced computing testbeds and supporting resources. The OLCF played a key role in successful completion of the ECP, including providing early access to exascale resources by industry and interagency partners, and supported multiple teams vying for ACM Gordon Bell Awards in 2022 and 2023.

The FY 2025 Request for this activity also supports operation and competitive allocation of the ALCF at ANL, including the Nation's second exascale system, an Intel/HPE-Cray system (Aurora) deployed in 2023 Funding also supports the 44 PF HPE/AMD/NVIDIA testbed (Polaris); an AI testbed program; IRI efforts, and supporting resources. ALCF's leadership in AI-enabled Exascale science was recognized by the 2022 ACM Gordon Bell Special Prize for HPC-based COVID-19 Research.

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 2024 INCITE allocations at the LCFs outpaced the available resources by a factor of three, 2023–2024 ALCC demand outpaced resources by a factor of five, and demand is expected to increase as ECP industry and interagency partners adopt ECP technologies as well as growing demand for AI resources. In addition, the LCFs play a key role in deploying IRI, with distinct challenges and resource requirements. Therefore, the LCFs have begun planning for upgrades that would expand capacity and capabilities to address both growing demand and expanded use cases. Meeting these challenges requires significant engagement with vendors to ensure availability of technologies that meet DOE mission needs. In FY 2025, 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 IRI.

High Performance Network Facilities and Testbeds
The FY 2025 Request for 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. The current generation of the ESnet backbone network, ESnet6, provides a new era of data transport orchestration, automation, and programmability that is foundational to DOE's IRI. ESnet is the circulatory system that enables the DOE science mission. The ESnet backbone network spans the continental U.S. and the Atlantic Ocean, connecting all 17 DOE National Laboratories and dozens of DOE sites to 200+ research and commercial networks around the world, enabling many tens of thousands of scientists across the country to access data and research resources. ESnet supports the data transport needs of all SC user facilities. In FY 2025, ESnet will continue to invest in site resiliency improvements across the DOE complex and will leverage ESnet6 to develop advanced services to support DOE priority R&D thrusts, DOE's IRI, and cybersecurity.

Integrated Research Infrastructure
This activity supports the community governance and operations of DOE's IRI. In FY 2025 IRI Operations will commence with seating of the IRI Management Council and initial investments to build core IRI services.

Advanced Scientific Computing Research
High Performance Computing and Network Facilities

Activities and Explanation of Changes

(dollars in thousands)

FY 2023 Enacted		FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
High Performance Computing and Network Facilities	**$652,003**	**$718,235**	**+$66,232**
High Performance Production Computing	$132,003	$146,500	+$14,497
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 early implementation of DOE's IRI.	The increase will support site preparations, design and procurements for the NERSC-10 upgrade, early implementation efforts for DOE's IRI, and sustaining ECP software and technologies critical to HCP operations and users.
National Energy Research Scientific Computing Center (NERSC)	*$130,000*	*$146,500*	*+$16,500*
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 the NERSC-9 Perlmutter system. In addition, funding will also support early implementation of DOE's IRI.	The increase will support site preparations, design, and long-lead procurement for the NERSC-10 upgrade, early implementation efforts for DOE's IRI, and sustaining ECP software and technologies critical to HPC operations and users.

(dollars in thousands)

FY 2023 Enacted	FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
High Performance Data Facility, OPC *$2,003*	*$ —*	*-$2,003*
Funding supports planning and preconceptual R&D for the HPDF, including site selection and preliminary design activities.	The Request reflects the planned advancement of HPDF as a line item project in FY 2024, contingent on achievement of CD-1 in FY 2024.	No OPC funding is requested.
Leadership Computing Facilities $430,000	$475,195	+$45,195
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.	Funding will support increased operating costs at both OLCF and ALCF to support allocation of the exascale systems, and payment of the Aurora system lease. 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.
Leadership Computing Facility at ANL *$175,000*	*$215,195*	*+$40,195*
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.	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. The Request will also support continuing operation and competitive allocation of the ALCF systems as well as AI and advanced computing testbeds, planning for future upgrades, vendor partnerships, and early implementation of DOE's IRI.	Funding will support 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 early implementation efforts for 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 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Leadership Computing Facility at ORNL *$255,000*	*$260,000*	*+$5,000*
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.	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 AI and advanced computing testbeds. OLCF will deploy and maintain ECP software and technologies critical to HPC operations and users. Summit will be decommissioned at the end of calendar year 2024. Planning for OLCF-6 will begin, including vendor engagements. The Request also supports early implementation of DOE's IRI.	Funding will support operating costs for the Frontier exascale system. Also, funding will support planning for future upgrades, vendor partnerships, and early implementation efforts for DOE's IRI, as well as maintenance of ECP software and technologies critical to HPC operations and users.
High Performance Network Facilities and Testbeds $90,000	$93,540	+$3,540
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 site resiliency investments and early implementation of DOE's IRI.	The increase will support operations of ESnet and early implementation efforts for DOE's IRI.
Integrated Research Infrastructure $ —	$3,000	+$3,000
	The Request will support commencement of IRI community governance activities and initial investments to build core IRI services.	The increase will support initiation of IRI governance and early operations investments.

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
Construction

Description

SC and NNSA completed the Exascale Computing Project (ECP), which was 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 all project KPPs, documentation, and close out activities. With the completion of the ECP, as well as upgrades at many of the other SC scientific user facilities, the Department recognized the need for a High Performance Data Facility (HPDF) to ensure the data from these scientific instruments would be accessible to the scientific community. In October 2023, ASCR determined site selection for 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 2025 Request includes $16,000,000 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. The project received approval for CD-0, Approve Mission Need, 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.

The HPDF will serve as a foundational element in enabling the DOE Integrated Research Infrastructure. HPDF 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 and data ecosystem for a wide variety of applications. DOE envisions HPDF will have a “Hub-and-Spoke” model in which a Hub will host centralized resources and also enable high priority DOE mission applications at “Spoke” sites by deploying and orchestrating distributed infrastructure at the Spokes or other locations. 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.

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.

Advanced Scientific Computing Research
Construction

Activities and Explanation of Changes

(dollars in thousands)

FY 2023 Enacted		FY 2025 Request	Explanation of Changes FY 2025 Request vs FY 2023 Enacted
Construction	**$77,000**	**$16,000**	**-$61,000**
17-SC-20, SC Exascale Computing Project	$77,000	$ —	-$77,000
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.		FY 2024 was the last year of funding for the project, which was successfully completed in FY 2024.	FY 2024 was the final year of funding for the ECP.
24-SC-20, High Performance Data Facility	$ —	$16,000	+$16,000
No funding was appropriated in FY 2023 for this project.		The Request will support design and early Spokes partnerships for the HPDF project in preparation for CD-2.	Funding will support design activities and early Spokes partnerships.

Advanced Scientific Computing Research
Capital Summary

(dollars in thousands)

	Total	Prior Years	FY 2023 Enacted	FY 2024 Annualized CR	FY 2025 Request	FY 2025 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 2023 Enacted	FY 2024 Annualized CR	FY 2025 Request	FY 2025 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) > $10M 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 2023 Enacted	FY 2024 Annualized CR	FY 2025 Request	FY 2025 Request vs FY 2023 Enacted
24-SC-20, High Performance Data Facility						
Total Estimated Cost (TEC)	293,000	-	-	-	16,000	+16,000
Other Project Cost (OPC)	8,076	2,000	2,003	4,000	-	-2,003
Total Project Cost (TPC)	**301,076**	**2,000**	**2,003**	**4,000**	**16,000**	**+13,997**
Total, Construction						
Total Estimated Cost (TEC)	**N/A**	**N/A**	**-**	**-**	**16,000**	**+16,000**
Other Project Cost (OPC)	**N/A**	**N/A**	**2,003**	**4,000**	**-**	**-2,003**
Total Project Cost (TPC)	**N/A**	**N/A**	**2,003**	**4,000**	**16,000**	**+13,997**

Note:

- *The current estimated TPC for the High Performance Data Facility is $304,933,000. In FY 2023, $1,930,000 in OPC funding was executed and is not reflected in this table.*

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 2023 Enacted	FY 2023 Current	FY 2024 Annualized CR	FY 2025 Request	FY 2025 Request vs FY 2023 Enacted
Scientific User Facilities - Type A					
National Energy Research Scientific Computing Center	**130,000**	**130,000**	**132,000**	**146,500**	**+16,500**
Number of Users	9,200	10,278	10,500	11,000	+1,800
Achieved Operating Hours	–	8,516	–	–	–
Planned Operating Hours	8,585	8,585	8,585	8,585	–
Unscheduled Down Time Hours	–	69	–	–	–
Argonne Leadership Computing Facility	**175,000**	**175,000**	**200,000**	**215,195**	**+40,195**
Number of Users	1,600	1,624	1,650	1,700	+100
Achieved Operating Hours	–	6,909	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Unscheduled Down Time Hours	–	99	–	–	–
Oak Ridge Leadership Computing Facility	**255,000**	**255,000**	**241,000**	**260,000**	**+5,000**
Number of Users	1,700	1,744	1,750	1,800	+100
Achieved Operating Hours	–	6,965	–	–	–
Planned Operating Hours	7,008	7,008	7,008	7,008	–
Unscheduled Down Time Hours	–	43	–	–	–
Energy Sciences Network	**90,000**	**90,000**	**88,400**	**93,540**	**+3,540**
Achieved Operating Hours	–	8,760	–	–	–
Planned Operating Hours	8,760	8,760	8,760	8,760	–
Scientific User Facilities - Type B					

	FY 2023 Enacted	FY 2023 Current	FY 2024 Annualized CR	FY 2025 Request	FY 2025 Request vs FY 2023 Enacted
	(dollars in thousands)				
High Performance Data Facility	**–**	**–**	**–**	**3,000**	**+3,000**
Total, Facilities	**650,000**	**650,000**	**661,400**	**718,235**	**+68,235**
Number of Users	12,500	13,646	13,900	14,500	+2,000
Achieved Operating Hours	–	31,150	–	–	–
Planned Operating Hours	31,361	31,361	31,361	31,361	–
Unscheduled Down Time Hours	–	211	–	–	–

Note:
- *Achieved Operating Hours and Unscheduled Downtime Hours will only be reflected in the Congressional budget cycle which provides actuals.*
- *Percent optimal operations defines what is achieved at this funding level. This includes staffing, up-to-date equipment and software, operations and maintenance, and appropriate investments to maintain world leadership.*

Advanced Scientific Computing Research
Scientific Employment

	FY 2023 Enacted	FY 2024 Annualized CR	FY 2025 Request	FY 2025 Request vs FY 2023 Enacted
Number of Permanent Ph.Ds (FTEs)	825	813	815	-10
Number of Postdoctoral Associates (FTEs)	365	341	345	-20
Number of Graduate Students (FTEs)	535	595	550	+15
Number of Other Scientific Employment (FTEs)	220	182	220	–
Total Scientific Employment (FTEs)	**1,945**	**1,931**	**1,930**	**-15**

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

24-SC-20, High Performance Data Facility
Thomas Jefferson National Accelerator Facility, TJNAF, and Lawrence Berkeley National Laboratory, LBNL
Project is for Design and Construction

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

Summary

The FY 2025 Request for the Office of Science (SC) High Performance Data Facility (HPDF) project is $16,000,000 of Total Estimated Cost (TEC) and $ — 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.

In October 2023, the Department announced the selection of the HPDF hub, which will create a new scientific user facility specializing in advanced infrastructure for data-intensive science. The Thomas Jefferson National Accelerator Facility (Jefferson Lab) will be the HPDF Hub Director and the lead infrastructure will be located at Jefferson Lab. The project to build the Hub will be a partnership between Jefferson Lab and Lawrence Berkeley National Laboratory (LBNL), and the two labs will form a joint project team led by Jefferson Lab charged to create an integrated HPDF Hub design.

HPDF will serve as a foundational element in enabling the DOE Integrated Research Infrastructure (IRI) and will provide crucial resources to SC and DOE programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources. HPDF 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 and data ecosystem for a wide variety of applications. DOE envisions HPDF will have a "Hub-and-Spoke" model in which a Hub will host centralized resources and also enable high priority DOE mission applications at "Spoke" sites by deploying and orchestrating distributed infrastructure at the Spokes or other locations. 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

The project was a new start in the FY 2024 Request. The most recent 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 Hub was selected via a merit review process. The FY 2024 Request will support conceptual design for the HPDF Hub project, an analysis of alternatives in preparation for CD-1, and potentially commencement of site preparation, contingent on achievement of CD-1 in FY 2024. The FY 2025 Request will support planning the final design and establishing the cost, scope, and schedule in preparation for CD-2/3 in late FY 2025.

Critical Milestone History

Fiscal Year	CD-0	Conceptual Design Complete	CD-1	CD-2	Final Design Complete	CD-3	CD-4
FY 2025	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
FY 2025	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 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 a “Hub and Spoke” model that will integrate centralized Hub and distributed Spoke resources to address mission essential streaming data and edge applications as a critical enabler of DOE’s IRI.

The Hub infrastructure is characterized by high-availability, high-performance data-centric resources designed with geographically and operationally resilient active-active failover, with the lead infrastructure at Jefferson Lab and resilience infrastructure at LBNL. The Spokes infrastructure will be distributed data-centric infrastructure at or near the edge, for example at SC User Facilities, DOE national laboratories, and US research institutions. The Hub and Spokes will be tied together as a single, integrated facility through orchestration hardware, software, and services operated by the Hub lead institution, Jefferson Lab, and its partner, LBNL.

The project scope will comprise: design, acquisition, delivery, and commissioning of the Hub infrastructure at Jefferson Lab and LBNL; design, acquisition, delivery, and commissioning of a set of initial spokes; integration of HPDF infrastructure with ESnet and the ASCR HPC facilities including NERSC, ALCF, and OLCF; software development for core HPDF services and development of an operations team that will support the infrastructure and scientific users; data center site preparation, power, and cooling infrastructure at Jefferson Lab and LBNL.

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 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 IRI 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 IRI vision.

The proposed HPDF will serve as a foundational element in enabling the DOE IRI; will provide crucial resources 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; 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)
In accordance with DOE Order 413.3B, the project will define preliminary KPPs at CD-1 and final KPPs at CD-2. The Threshold KPPs will represent the minimum acceptable performance that the project must achieve. Achievement of the Threshold KPPs will be a prerequisite for approval of CD-4, Project Completion. The Objective KPPs will represent the desired project performance.

Performance Measure	Threshold	Objective
Design/construct building	TBD	TBD
Instrumentation design/development	TBD	TBD

3. Financial Schedule

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Estimated Cost (TEC)			
Design (TEC)			
FY 2024	1,000	1,000	—
FY 2025	3,000	3,000	3,500
Outyears	—	—	500
Total, Design (TEC)	**4,000**	**4,000**	**4,000**
Construction (TEC)			
FY 2025	13,000	13,000	13,000
Outyears	277,000	277,000	277,000
Total, Construction (TEC)	**290,000**	**290,000**	**290,000**
Total Estimated Cost (TEC)			
FY 2024	1,000	1,000	—
FY 2025	16,000	16,000	16,500
Outyears	277,000	277,000	277,500
Total, Total Estimated Cost (TEC)	**294,000**	**294,000**	**294,000**

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Other Project Cost (OPC)			
Prior Years	1,930	1,930	–
FY 2023	1,930	1,930	–
FY 2024	7,000	7,000	7,933
FY 2025	–	–	2,927
Outyears	73	73	73
Total, Other Project Cost (OPC)	**10,933**	**10,933**	**10,933**

(dollars in thousands)

	Budget Authority (Appropriations)	Obligations	Costs
Total Project Cost (TPC)			
Prior Years	1,930	1,930	–
FY 2023	1,930	1,930	–
FY 2024	8,000	8,000	7,933
FY 2025	16,000	16,000	19,427
Outyears	277,073	277,073	277,573
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	2,600	**N/A**
Design - Contingency	1,400	1,400	**N/A**
Total, Design (TEC)	**4,000**	**4,000**	**N/A**
Construction	188,500	188,500	**N/A**
Construction - Contingency	101,500	101,500	**N/A**
Total, Construction (TEC)	**290,000**	**290,000**	**N/A**
Total, TEC	**294,000**	**294,000**	**N/A**
Contingency, TEC	*102,900*	*102,900*	***N/A***
Other Project Cost (OPC)			
OPC, Except D&D	7,106	7,106	**N/A**
OPC - Contingency	3,827	3,827	**N/A**

(dollars in thousands)

	Current Total Estimate	Previous Total Estimate	Original Validated Baseline
Total, Except D&D (OPC)	**10,933**	**10,933**	**N/A**
Total, OPC	**10,933**	**10,933**	**N/A**
Contingency, OPC	*3,827*	*3,827*	***N/A***
Total, TPC	**304,933**	**304,933**	**N/A**
Total, Contingency (TEC+OPC)	***106,727***	***106,727***	***N/A***

5. Schedule of Appropriations Requests

(dollars in thousands)

Fiscal Year	Type	Prior Years	FY 2023	FY 2024	FY 2025	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
FY 2025	TEC	—	—	1,000	16,000	277,000	294,000
	OPC	1,930	1,930	7,000	—	73	10,933
	TPC	1,930	1,930	8,000	16,000	277,073	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	TBD

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

In conjunction with the HPDF project, Jefferson Lab will design, construct, and commission the Jefferson Laboratory Data Center (JLDC) building using Commonwealth of Virginia funds. The JLDC scope will contain construction of the building and the infrastructure work required for the HPDF project to take beneficial occupancy on the schedule necessitated by the HPDF project.