2023 Advanced Scientific Computing Advisory Committee Facilities Charge

In response to Dr. Berhe's December 2023 Facilities Charge, Advanced Scientific Computing Research (ASCR) is providing the following list of projects for consideration by the subcommittee.

Leadership Computing Program

Argonne Leadership Computing Facility 4 (ALCF-4)
Total Project Cost (TPC) range at CD-0: $200M to $328M
HPC System Acquisition cost range at CD-0: $200M to $350M
Critical Decision 0 approved May, 2023 (note: ALCF-4 and OLCF-6 have a joint CD-0)
This project responds to the statutory requirement that DOE provide leadership class computing resources on a competitive, merit-review basis to researchers in U.S. industry, institutions of higher education, National Laboratories, and other Federal agencies. DOE ASCR's Leadership Computing program manages risk through two facilities at Argonne and Oak Ridge National Laboratories that provide hardware architectural diversity and staggered deployment timelines. The ALCF-3 (Aurora) supercomputing system life-cycle is expected to be mid-2024 through mid-2029, with possible life extension to mid-2030; the ALCF-4 project will therefore target completion in 2029. ALCF-4 will continue to provide DOE and the nation with access to leadership-class computing---the world's most advanced extreme-scale supercomputing resources, with emphasis on advancing US competitiveness in energy efficient computing and AI. The project scope encompasses the site preparation, power, and cooling necessary for the ALCF-4 system; public-private R&D partnership(s) with vendors for new computing hardware; and system installation and commissioning. The ALCF-4 project will: partner with US vendors to drive innovation in energy efficient AI-accelerated hardware; stake out a leadership position for U.S. science in advancing AI workflows and large AI model development for priority DOE science applications; deploy Integrated Research Infrastructure services and infrastructure in collaboration with HPDF, ESnet, NERSC, and OLCF, and other SC User Facilities; and refresh and expand data storage systems and analysis clusters to accommodate massive-scale AI projects.

Argonne Leadership Computing Facility 5 (ALCF-5)
Total Project Cost (TPC) range: TBD
The ALCF-4 supercomputing system life-cycle is expected to be mid-2029 through mid-2034, with possible life extension to mid-2035; the ALCF-5 project will therefore target completion in 2034. The ALCF-5 project will provide the next generation of leadership class computing resources and will forge vendor partnerships for U.S. competitive advantage in energy-efficient computing, AI accelerators, and hybrid-quantum computing.

Oak Ridge Leadership Computing Facility 6 (OLCF-6)
Total Project Cost (TPC) range at CD-0: $200M to $328M
HPC System Acquisition cost range at CD-0: $200M to $350M
Critical Decision 0 approved May, 2023 (note: ALCF-4 and OLCF-6 have a joint CD-0)
This project responds to the statutory requirement that DOE provide leadership class computing resources on a competitive, merit-review basis to researchers in U.S. industry, institutions of higher education, National Laboratories, and other Federal agencies. DOE ASCR's Leadership Computing program manages risk through two facilities at Argonne and Oak Ridge National Laboratories that provide hardware architectural diversity and staggered deployment timelines. The OLCF-5 (Frontier) supercomputing system life-cycle is expected to be mid-2023 through mid-2028, with possible life extension to mid-2029; the OLCF-6 project will therefore target completion in 2028. OLCF-6 will continue to provide DOE and the nation with access to leadership-class computing---the world's most advanced extreme-scale supercomputing resources, with emphasis on advancing US competitiveness in energy efficient computing and AI. The project scope encompasses the site preparation, power, and cooling necessary for the OLCF-6 system; public-private R&D partnership(s) with vendors for new computing hardware; and system installation and commissioning. The OLCF-6 project will: partner with US vendors to achieve more energy efficient AI-accelerated hardware; stake out a leadership position for U.S. science in hybrid quantum/HPC computing for priority DOE science applications; provide an enhanced FISMA-moderate computing/data enclave, which provides a unique platform for AI research involving protected or sensitive data; deploy Integrated Research Infrastructure services and infrastructure in collaboration with HPDF, ESnet, NERSC, and OLCF, and other SC User Facilities; and refresh and expand data storage systems and analysis clusters to accommodate massive-scale AI projects.

Oak Ridge Leadership Computing Facility 7 (OLCF-7)
Total Project Cost (TPC) range: TBD
The OLCF-6 supercomputing system life-cycle is expected to be mid-2028 through mid-2033, with possible life extension to mid-2034; the OLCF-7 project will therefore target completion in 2033. The OLCF-7 project will provide the next generation of leadership class computing resources, and will forge vendor partnerships for U.S. competitive advantage in energy-efficient computing, AI accelerators, and hybrid-quantum computing.

High Performance Production Computing Program

National Energy Research Scientific Computing Center 10 (NERSC-10)
Total Project Cost (TPC) range at CD-0: $80M to $185M
Note: The project is proposing a revision of the TPC to a range of $136M to $224M at CD-1
HPC System Acquisition cost range: $150M to $400M
Critical Decision 0 approved October 2021
This project responds to Office of Science mission requirements to provide the Office of Science research community, including SC User Facilities, with dedicated access to large-scale high performance computing resources. The NERSC-9 (Perlmutter) supercomputing system life-cycle is expected to be mid-2022 through mid-2027, with possible life extension to mid-2028; the NERSC-10 project will therefore target completion in 2027. The HPC requirements of this broadening user community is increasingly diverse, with a growing demand for complex high performance workflow capabilities from our user community across all scientific domains. HPC workflows interconnect computational and data manipulation steps, and require significant computational, storage and/or network resources within. The NERSC-10 system will integrate available new technologies and support the emerging needs in AI and experimental/observational science to accelerate end-to-end DOE SC workflows and enable new modes of scientific discovery through the integration of experiment, data analysis, and simulation. The project simultaneously brings the scientific impact of the de-risked hardware-software ecosystem achieved through the Exascale Computing Initiative to NERSC's vast user community while creating a malleable high performance computing platform for the Integrated Research Infrastructure ecosystem. The NERSC-10 project will support the emerging needs in AI and experimental/observational science by accelerating end-to-end DOE SC workflows and enabling new modes of scientific discovery through the integration of experiment, data analysis, and simulation. The project scope encompasses the site preparation, power, and cooling necessary for the NERSC-10 system; public-private R&D partnership(s) with vendors for new computing hardware; and system installation and commissioning. NERSC-10 will deliver large-scale hybrid High Performance Computing resources, capable of running hundreds of diverse workflows simultaneously, to the computational science community and advance Integrated Research Infrastructure services for complex multi-facility integrated science, AI-driven workflows, and real-time interactive computing in collaboration with HPDF, ESnet, ALCF, and OLCF, and other SC User Facilities.

National Energy Research Scientific Computing Center 11 (NERSC-11)
Total Project Cost (TPC) range: TBD
The NERSC-10 supercomputing system life-cycle is expected to be mid-2027 through mid-2032, with possible life extension to mid-2033; the NERSC-11 project will therefore target completion in 2032. The NERSC-11 project will provide the next generation of HPC resources to the SC research community, coupling closely with ESnet and HPDF to advance the IRI ecosystem.

High Performance Data Facility (HPDF)
Total Project Cost (TPC) range at CD-0: $300M to $500M
Critical Decision 0 approved October 2020

This project responds to Office of Science mission requirements to provide the Office of Science and DOE research community, including SC User Facilities and the DOE National Laboratories, with dedicated access to high performance data management infrastructure for data-intensive research. The project scope is to design and build a new scientific user facility that fits within and adds world-class capabilities to the ASCR and SC data and computing infrastructure ecosystem. The facility's mission will be to enable and accelerate scientific discovery by delivering state-of-the-art data management infrastructure, capabilities, and tools. HPDF is envisioned as a hub-and-spoke model, in which the Hub will both host centralized resources and also enable high-priority DOE mission applications at spoke sites by deploying, managing, and orchestrating distributed infrastructure at the spokes. The number and variety of spokes is likely to grow and evolve with mission requirements. On October 16, 2023 DOE announced the meritorious selection of Thomas Jefferson National Accelerator Facility (Jefferson Lab) as the HPDF Hub Director and sited the lead infrastructure on the lab's campus in Newport News, Virginia, and determined that the HPDF project will be a partnership between Jefferson Lab and Lawrence Berkeley National Laboratory (LBNL). HPDF will be a first-of-its-kind Office of Science (SC) user facility A successful, fully-realized HPDF will be widely recognized as a national and international leader in uplifting data science and high performance data infrastructure. The project scope will comprise the following: design, acquisition, delivery and commissioning of the centralized Hub hardware at Jefferson Lab and LBNL; design, acquisition, delivery and commissioning of initial Spokes hardware; network integration of HPDF infrastructure with ESnet; software development for core HPDF services; data center site preparation, power, and cooling infrastructure at Jefferson Lab and LBNL. 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.

High Performance Data Facility (HPDF) Spokes 1 and Spokes 2
Total Project Cost (TPC) range: TBD
This project responds to Office of Science mission requirements to provide the Office of Science and DOE research community, including SC User Facilities and the DOE National Laboratories, with dedicated access to high performance data management infrastructure for data-intensive research. The project scope is to extend the number of installed HPDF spokes, extending the reach and impact of HPDF and the ASCR computing and data infrastructure ecosystem to even more institutions, user facilities, and/or partners. The Spokes 1 and Spokes 2 projects will occur after the HPDF project has completed (i.e., after CD-4) and will deploy distributed infrastructure at the spokes.

High Performance Networking Program

Energy Sciences Network 7 (ESnet7)
Total Project Cost (TPC) range: TBD
This project responds to Office of Science mission requirements to provide high reliability (> 99.9% uptime on per annum operations baseline) high performance wide-area network capabilities, tailored for a diverse range of research, programmatic, and institutional demands, to the Department of Energy and other critical stakeholders. The project scope is to refresh the optical, routing, and switching hardware that constitute the core of the ESnet backbone network for enhanced performance and cyber security; add installed network traffic capacity to meet growth demands; extend Integrated Research Infrastructure capabilities through deployed hardware and advanced data transport services, including such concepts as: In-Network Caching, Data Transfer Nodes as-a-service, SmartEdge Gateway, and integration of SENSE capabilities for automated workflow optimization; and extend the reach of DOE research infrastructure to new institutional partners. The ESnet6 project completed in 2022; the refresh cycle for core optical, routing, and switching hardware will necessitate deployment of ESnet7 in the 2027-2029 timeframe.

Energy Sciences Network 8 (ESnet8)
Total Project Cost (TPC) range: TBD
This project responds to Office of Science mission requirements to provide high reliability (> 99.9% uptime on per annum operations baseline) high performance wide-area network capabilities, tailored for a diverse range of research, programmatic, and institutional demands, to the Department of Energy and other critical stakeholders. Assuming the ESnet7 project is completed in the 2027-2029 timeframe; the refresh cycle for core optical, routing, and switching hardware will necessitate deployment of ESnet8 in the 2032-2035 timeframe.