

LEADERSHIP-CLASS COMPUTING FACILITY (LCCF)**\$201,000,000****Requested MREFC Funds¹**

(Dollars in Millions)

	FY 2026 Request
Leadership-Class Computing Facility	\$201.00

¹ The previously presented Total Project Cost of \$457.45 million for LCCF remains intact.

Brief Description

Computer simulation, together with artificial intelligence (AI) methods and data analytics, is critical to enabling transformational science and engineering (S&E) research. From understanding the origin and evolution of our universe to exploring atomic-scale biomolecular processes, computational methods are an integral part of almost all curiosity-driven, use-inspired, and translational S&E research. LCCF, led by the Texas Advanced Computing Center (TACC) at the University of Texas at Austin, is envisioned as a distributed facility that will provide unique computational and data analytics capabilities, as well as critical software and services, for the nation's S&E research community to enable discoveries that would not be possible otherwise. Furthermore, the project will deploy a broad portfolio of education and outreach activities that will expand and nurture our nation's future S&E workforce in data and computational sciences.

LCCF's core processing capabilities will be anchored by a computing system called *Horizon*, which will represent a substantial increase over the computational and data analytics capacity of the current NSF leadership-class computing system, *Frontera*. LCCF will also include four Distributed Science Centers (DSCs), which will be located at the Atlanta University Center Consortium (AUCC), the National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign, the Pittsburgh Supercomputing Center (PSC), and the San Diego Supercomputer Center (SDSC). These DSCs are designed to leverage expertise across the broader cyberinfrastructure ecosystem and to provide critical edge-computing services to the LCCF user community. The partnership with AUCC will leverage the consortium's Data Science Initiative to provide workforce pathways into leadership computing. The partnerships with NCSA and PSC will explore and provide new processor technologies for AI, and data intensive computing and data mirrors for published archives, respectively. Finally, the partnership with SDSC will focus on supporting Machine Learning (ML) and instrument data analytics in scientific workflows, and methods to democratize access to LCCF.

In addition to supporting large-scale simulations and AI, LCCF will also facilitate new usage modes such as interactive computing required by scientific tasks that involve human-in-the-loop processing, as well as urgent computing for emergency response scenarios that will need immediate access to computing resources and real-time data. LCCF will support the full scientific data lifecycle, which is critical to modern S&E discovery processes and to unleashing the potential of rapid advancements in ML and AI. The extensive LCCF Education and Public Outreach (EPO) effort will expand the learning and workforce pipeline and aims to attract talented individuals to research computing. A key

component of the LCCF EPO activities will be a Visitor Center at TACC that will include virtual and physical exhibits, along with multiple pilot activities that will be designed to engage students, teachers, and the broader public across the nation.

Baseline History

In Fiscal Year 2017, NSF released solicitation NSF 17-558 (*Towards a Leadership-Class Computing Facility - Phase 1*)¹ to support the acquisition of a Phase 1 leadership-class computing system and to initiate the planning process for a future Phase 2 LCCF. After rigorous merit review of the submitted proposals, NSF made an award to TACC for the acquisition of the *Frontera* system, the first acquisition in the two-phase process, and to advance the planning for the Phase 2 LCCF. As noted in NSF 17-558, and in response to the recommendations set forth in the report *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020*,² the goal of the LCCF Phase 2 planning was the design of a major facility that would support all S&E research.

The Major Facilities Design Stage, as defined in the NSF Research Infrastructure Guide,³ consists of three phases – Conceptual Design, Preliminary Design, and Final Design. In July 2019, LCCF began the Conceptual Design Phase following formal admission to the Design Stage by the NSF Director. The Conceptual Design Review (CDR) was conducted in June 2020 by an external panel of experts to evaluate the proposed, site-independent design of LCCF. A successful CDR resulted in the NSF Director's approval for the project to enter the Preliminary Design Phase in September 2020.

Following extensive Preliminary Design Phase planning activities, the project successfully completed the Preliminary Design Review (PDR) in January 2022. At PDR, the project was judged by an external panel of experts to have made appropriate progress in clearly articulating a cost estimate and a near-final project definition, as well as completing a sufficiently mature risk analysis to allow determination of the risk-adjusted Total Project Cost (TPC) and construction duration for establishing a budget request to Congress. As a result of the successful PDR, the project was approved by the NSF Director to enter the Final Design Phase in August 2022.

A Final Design Review (FDR) of the project was conducted in April 2023. An expert external review panel evaluated the readiness of the project to enter the Construction Stage and recommended that LCCF be advanced. At its November 2023 meeting, the National Science Board discussed LCCF and passed a resolution in favor of advancing the project to the construction stage. In December 2023, the NSF Director authorized the advancement of the project to the construction stage. A construction award was made for the project in July 2024.

Project Status

In August 2024, the project signed a colocation agreement with Sabey Inc. to provide 32,000 square feet of customized datacenter space for high-density utility hosting of the primary LCCF computing system called *Horizon*. Since entering into the agreement, the project has been holding weekly

¹ www.nsf.gov/pubs/2017/nsf17558/nsf17558.htm

² www.nationalacademies.org/our-work/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-in-2017-2020

³ www.nsf.gov/pubs/2021/nsf21107/nsf21107.pdf

meetings with the Sabey construction team to ensure on-time delivery of the datacenter colocation space, expected in March 2026.

In early 2025, the project team continued with preparations for making other initial procurements, including the *Horizon* system, visitor center construction, networking equipment, and the prototype system and software that will be available for early-access user testing of the LCCF computing environment. The Project continues to hold collaborative technology discussions with Dell, NVIDIA, and other partners to closely monitor future technological roadmaps. No outstanding technical issues remain.

Meeting Intellectual Community Needs

The scientific requirements for LCCF are defined by the needs of the S&E community. These include the need to expand computation and data analytics capabilities, as well as the size and breadth of the workforce that will use them. LCCF requirements also include agility and the ability to adapt to rapidly changing technology, new application formulations and requirements, as well as new paradigms of computation-based research. Furthermore, LCCF will need to enhance the broader high-performance computing ecosystem by providing unique resources and expertise at our nation's Institutions of Higher Education that complement other investments made by the federal government and in industry.

The LCCF scientific requirements were determined through a broad set of engagements with the S&E community. These included: distillation from a wide variety of scientific community reports; analysis of usage patterns on leadership computing systems across government; feedback from project-organized planning meetings with the research community;⁴ discussions under non-disclosure agreements with industry vendors concerning technology roadmaps; technical evaluation of advanced system prototypes; and direct conversations and interviews with scientists and engineers.⁵

During the Design Stage, LCCF augmented its gathering of science requirements by selecting twenty-one Characteristic Science Application (CSA) teams as project partners. The CSA partners represent applications across a broad range of S&E domains and were selected to enable the development of a suite of science drivers that will be used to verify and validate the facility's effectiveness. The CSA partners also provided requirements to inform design decisions and will continue to enable the project to acquire experience and expertise in coding/performance-tuning enhancements in preparation for facility operations.

Finally, LCCF will be a key enabler of large-scale AI research to ensure continued U.S. world leadership in S&E research and development. In particular, the facility will be a key element in advancing a national research infrastructure to broaden access for the entire nation to critically-needed AI resources, as envisioned in the White House *National AI Research Resource Task Force* report.⁶ When

⁴ lccf.tacc.utexas.edu/media/filer_public/d4/e4/d4e49b4f-26b2-4065-995e-981f9e21b03f/lccf-tec-science-requirement_workshop.pdf

⁵ lccf.tacc.utexas.edu/media/filer_public/24/c8/24c85369-f2ea-4e57-9647-be7ad656e361/lccf-req001_sc19_bof_report.pdf

⁶ www.ai.gov/nairrtf/

operational, LCCF will advance U.S. dominance in AI, as envisioned by the President's Executive Order on *Removing Barriers to American Leadership in Artificial Intelligence*.⁷

Governance Structure and Partnerships

NSF Governance Structure

NSF oversight of LCCF is provided by a Program Officer in the CISE Office of Advanced Cyberinfrastructure (OAC), working cooperatively with other NSF staff through an Integrated Project Team (IPT). The LCCF IPT includes members from OAC leadership, Research Infrastructure Office (RIO), Office of the General Counsel (OGC), Budget Division, and Division of Acquisition and Cooperative Support. In addition, RIO staff provides advice to OAC and assists with agency oversight and assurance. CISE leadership and NSF's Chief Officer for Research Facilities provide additional high-level guidance and oversight to the project.

External Governance Structure

The LCCF Senior Management Team consists of the LCCF Project Director (PD), the Project Manager (PM), and the Deputy Project Manager (DPM). This group is responsible for the day-to-day management of the LCCF project. The PD serves as the primary interface between the project and NSF's oversight team. The Project Management office is headed by the PM and DPM, who report to the PD. The PM has line responsibility for operations, reporting, and process management within the project, with support from the DPM. The PD is advised by two groups: the Technology Advisory Board and the Science Advisory Board. Each board consists of members who are leaders in their fields and represent cyberinfrastructure providers and experts (the Technology board) and stakeholders from the community of scientists who will be the ultimate customers of LCCF (the Science board).

Partnerships and Other Funding Sources

The LCCF project includes partnerships with more than twenty academic institutions that will contribute to design, validation, and eventual operations. The project also has extensive industry partnerships in various high-tech economic sectors, such as computing hardware, software, and data center colocation.

Cost and Schedule

The project began construction in July 2024. The project construction schedule has a duration of 34 months, providing for facility acceptance in FY 2027. The project has also developed schedule contingency of one year to allow for uncertainties in the construction process.

Future Operations Costs

Contingent on the availability of funds, NSF plans to support LCCF operations and maintenance for an initial five years with a possibility of renewal for another five years. The current estimate is that LCCF operations will cost approximately \$40 million annually. NSF support for LCCF beyond 10 years, and possible re-competition of the operations and maintenance award, will be informed by a planned future study by the National Academies of Sciences, Engineering, and Medicine or other similar

⁷ www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/

reports by community-based bodies, and will follow internal NSF procedures for renewal, competition, and disposition.

The current estimated cost for disposition of LCCF at the end of the ten-year operations period is \$0. With the eventual discontinuation of NSF support for LCCF, the main data center will be returned to the colocation provider, and the computing system will be repurposed by the recipient. There is no expectation that site restoration, demolition, or modification would be needed when the facility is decommissioned.

Reviews

Technical Reviews

The LCCF project was technically reviewed multiple times during the Design Stage by external committees of experts. These include stage gate reviews at CDR, PDR, and FDR in FYs 2020, 2022 and 2023, respectively. The LCCF CSA program was also reviewed by a panel of external experts in FY 2021.

Management, Cost, and Schedule Reviews

In accordance with NSF's No Cost Overrun Policy, the agency has developed a risk-adjusted TPC estimate post-PDR (based on known risks). Using the TACC proposal as the basis, BFA employed a series of independent cost estimate reviews to inform NSF's post-PDR cost analysis. These included an independent cost estimate in accordance with U.S. Government Accountability Office good practices, conducted by an external contractor for appropriate project components, by BFA's internal Cost Analysis and Pre-award Branch for other elements of cost, and by RIO for assessments of budget contingency and schedule. This information was reconciled against the original proposed budget from TACC in determining the recommended value for inclusion in the FY 2024 Budget Request.

At the completion of FDR, a refined final risk-adjusted TPC and schedule were proposed by the project. BFA conducted a second series of independent cost estimate reviews in accordance with NSF policies prior to the issuance of the Construction award.

Risks

Technical

The LCCF project deploys state-of-the-art technologies and services for the Nation's S&E research community. Several risks related to the readiness and reliability of the future deployed systems and services are included in the LCCF risk register. The LCCF CSA partnerships will be key in monitoring and managing these risks as the CSA teams will be the early users of the facility and will help to identify and remedy issues as they arise. Cost and schedule contingencies have also been factored into the plan to mitigate these risks. In addition, a scope management plan provides options to mitigate cost and/or schedule overruns by eliminating project scope if necessary. LCCF's scope management plan is arranged so that any decision to descope can be made as late as possible in the project and with minimal impact on the science goals.

Environmental Health and Safety

The project includes several health and safety-related risks in its risk register, including events that may have a system-wide impact on the LCCF construction schedule. Contingencies have been factored into the construction schedule to mitigate these risks if necessary. Furthermore, an assessment for

compliance with National Environmental Policy Act regulations has been completed by NSF OGC.

Partnership Risk

The project has established partnerships with several technology vendors. Risks of withdrawal by partners due to unforeseen business reasons are included in the risk register and are carefully monitored and managed by the project. Project assessment of these partnership risks is that the probability of their occurrence is very low. However, mitigation plans have been developed, including vendor penalties for late delivery.

System Integration Risk

An experienced technology integrator for the core *Horizon* system has been identified, and overall system integration risk is low. However, several other risks have been identified by the project related to delays in data center access and technology roadmap changes that could impact the system integration timeline and schedule. Cost and schedule contingencies have been factored into the plan to mitigate these risks.