DIRECTORATE FOR COMPUTER AND INFORMATION SCIENCE AND ENGINEERING (CISE)

$1,150,780,000
+$143,650,000 / 14.3%

CISE Funding
(Dollars in Millions)

	FY 2021[1] Actual	FY 2021 ARP Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Percent
Office of Advanced Cyberinfrastructure (OAC)	$230.44	$6.59	-	$252.25	$21.81	9.5%
Computing and Communication Foundations (CCF)	200.95	1.75	-	218.57	$17.62	8.8%
Computer and Network Systems (CNS)	238.02	4.87	-	266.06	$28.04	11.8%
Information and Intelligent Systems (IIS)	217.78	1.75	-	248.16	30.38	14.0%
Information Technology Research (ITR)	119.94	20.75	-	165.74	45.80	38.2%
Total	**$1,007.13**	**$35.72**	**-**	**$1,150.78**	**$143.65**	**14.3%**

[1] Funding for FY 2021 is adjusted for comparability to reflect the movement of I-Corps™ to TIP in FY 2022.

About CISE

Advances in information technology (IT) over recent decades have proven to be key drivers of the U.S. economy. Essentially all practical applications of today's IT are based on ideas and concepts that emerged from investments in fundamental computing and information research, many of them funded by CISE.[1] Fundamental ideas and concepts advanced through computing and information research have enabled innovative products and applications that now benefit many aspects of daily life, including personal communication, clean energy, intelligent transportation, health care, advanced manufacturing, national and homeland security, disaster preparedness and response, education and workforce development, public and private organizational effectiveness and efficiency, and discovery and innovation at the frontiers of all areas of scientific and engineering research. CISE investments will accelerate climate and clean energy research, advance equity in science, engineering, and society, and bolster U.S. leadership in critical and emerging technologies.

CISE's mission is to promote the progress of computer and information science and engineering research and education, and advance the development and use of cyberinfrastructure (CI) across the science and engineering research enterprise; to promote understanding of the principles and uses of advanced computer, communication, and information systems in advancing science and engineering and in service to society; and to contribute to universal, transparent, and affordable participation in a knowledge-based society. CISE supports ambitious research and research infrastructure projects within and across the many subfields of computing, as well as advanced research CI for all areas of science and engineering; contributes to the education and training of computing and information professionals; and more broadly, informs the preparation of a U.S. workforce with computing, computational, and information competencies essential for success in an increasingly competitive global and digital market. CISE investments foster and support research and teaching environments that promote equity. CISE executes its mission through its Divisions of Computing and Communication Foundations (CCF), Computer and Network Systems (CNS), Information and Intelligent Systems (IIS), and Information Technology Research (ITR), and through the Office of Advanced Cyberinfrastructure

[1] www.nap.edu/catalog/25961/information-technology-innovation-resurgence-confluence-and-continuing-impact

(OAC), which has a Foundation-wide role supporting advanced research CI for all areas of science and engineering—and in close partnership with other NSF units, federal agencies, private industry and foundations, and international funders.

In FY 2023, CISE will continue to support Nation's priorities through investments in AI, advanced computing systems and services including high-performance computing (HPC), QIS, advanced communications technologies, advanced manufacturing, semiconductors and microelectronics, biotechnology, cybersecurity, and disaster response and resilience. CISE's investments in these areas contribute significantly to national security, economic competitiveness, sustainability, and the broad advancement of all fields of science and engineering. Advances in these areas will provide opportunities for major scientific breakthroughs and will positively transform U.S. lives and industry for years to come. As part of these investments, CISE will advance democracy-affirming technologies, including privacy-preserving technologies.

CISE's FY 2023 Budget Request is also shaped by the directorate's continued support for NSF's Big Ideas, including co-leadership of HDR and FW-HTF, and participation in NNA and URoL. Further, as part of HDR, and in partnership with the other research directorates and offices, CISE will invest, through its ITR division, in convergent activities that transcend the traditional disciplinary boundaries of individual NSF units. CISE, as the steward for HDR and in partnership with the other directorates, will support fundamental research in data science and engineering; development of a cohesive, federated approach to the research data infrastructure; and development of a 21st-century data-capable workforce. CISE's FY 2023 Budget Request comprises support for other ongoing NSF-wide priorities as well, including microelectronics and semiconductor research and SaTC.

CISE, through OAC, will provide NSF's co-leadership of the National Science and Technology Council's Future Advanced Computing Ecosystem (FACE) subcommittee.[2] As part of its support for FACE, CISE investments will support the full breadth of NSF-funded S&E, including research furthering our understanding of climate science and clean-energy technologies, by (i) advancing future computing paradigms, devices, architectures, and platforms; and (ii) furthering the development and deployment of advanced computing systems and services, including maximizing the benefits of these systems and services through the deep integration of emerging computing paradigms with current science and engineering research drivers. These investments will enable shared resources and improved capabilities across a range of disciplines, a diverse set of users within a large number of academic institutions, and a wide range of science and engineering advances. In FY 2023, CISE will invest, through its ITR division, in the development of a National Discovery Cloud (NDC) for Climate. This resource will federate advanced compute, data, software and networking resources, democratizing access to a cyberinfrastructure ecosystem that is increasingly necessary to further climate-related S&E. The NDC for Climate will serve as a pilot for future efforts to enable equitable access to an NDC across all fields of S&E.

Given the increasingly influential societal role of computing research, it is critical to ensure the Nation offers broad access to and education on this topic. As a part of an agency-wide emphasis, CISE will continue to invest in a broad suite of activities to support broadening participation in research and education in CISE fields and STEM more generally. For example, in alignment with NSF INCLUDES, the Broadening Participation in Computing Alliances (BPC-A) will serve as broad coalitions of institutions

[2] www.nitrd.gov/pubs/Future-Advanced-Computing-Ecosystem-Strategic-Plan-Nov-2020.pdf

of higher education, K-12 schools, government, industry, professional societies, and other not-for-profit organizations that design and carry out comprehensive programs addressing underrepresentation in the computing and information science disciplines. Additionally, the CISE Minority-Serving Institutions Research Expansion (CISE-MSI) program will continue to broaden participation by increasing the number of CISE-funded research projects from MSIs, which are central to inclusive excellence. CISE's investments in Computer Science for All (CSforAll) and CISE Graduate Fellowships (CSGrad4US) also will emphasize education and training of more U.S.-based students from diverse backgrounds.

CISE will continue to provide leadership for the Federal Government's Networking and Information Technology Research and Development (NITRD) program. The NITRD Subcommittee of the National Science and Technology Council (NSTC), which coordinates investments in networking and information technology research and development across more than 20 federal departments, agencies, and offices, is co-chaired by the NSF assistant director for CISE. All research, education, and research infrastructure projects supported by CISE contribute to NSF's NITRD portfolio. In addition, CISE co-chairs the National Artificial Intelligence Research Resource (NAIRR) Task Force, which is charged with developing a roadmap and implementation plan for a shared computing and data infrastructure. The envisioned NAIRR aligns with the NDC for Climate described above.

Finally, CISE will build, strengthen, and expand strategic, multisector partnerships, including those with other NSF units, other federal agencies, private industry and foundations, and international funders, as an increasingly important means to maximize the scientific, economic, and societal impacts of the directorate's investments. These external partnerships leverage resources, inform use-inspired research, accelerate the translation of research innovations to practice, and enhance workforce development. CISE will coordinate closely with other directorates as well as with the Strategic Partnerships Office within the TIP directorate.

CISE provides about 79 percent of the federal funding for fundamental computer science research at U.S. academic institutions.

Major Investments

CISE Major Investments

(Dollars in Millions)

Area of Investment[1,2]	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Change over FY 2021 Actual Percent
Advanced Manufacturing	$44.40	-	$42.22	-$2.18	-4.9%
Advanced Wireless Research	87.45	-	93.26	5.81	6.6%
Artificial Intelligence	344.00	-	369.80	25.80	7.5%
Biotechnology	6.92	-	6.00	-0.92	-13.3%
CISE Graduate Fellowships	-	-	20.70	20.70	N/A
Climate: Clean Energy Technology	24.22	-	31.12	6.90	28.5%
Climate: USGCRP[3]	-	-	40.00	40.00	N/A
Microelectronics and Semiconductors	17.95	-	23.46	5.51	30.7%
Quantum Information Science	20.70	-	24.28	3.58	17.3%
Secure & Trustworthy Cyberspace[4]	70.81	-	75.81	5.00	7.1%

[1] Major investments may have funding overlap and thus should not be summed.

[2] This table reflects this directorate's support for selected areas of investment. In other directorate narratives, areas of investment displayed in this table may differ and thus should not be summed across narratives.

[3] Funding includes resources for agency-wide initiatives.

[4] FY 2023 Request funding includes $5.0 million to strengthen the national cybersecurity workforce pipeline, which is complementary to the Cyber Defense Education & Training program at the Cybersecurity and Infrastructure Security Agency. An additional $5.0 million for this program is budgeted within EDU.

- Advanced Manufacturing: CISE will invest in research that integrates ubiquitous sensors, computational tools, and highly connected cyber-physical systems in smart processing and "cyber-manufacturing" systems. This investment will enable new functionalities that will increase the efficiency and sustainability of the production of the next generation of products and services.
- Advanced Wireless Research: CISE will continue to invest in research in advanced wireless networks, building on its track record of enabling early-stage successes in 5G through ground-breaking millimeter-wave research. CISE investments will specifically enable further exploration of additional spectrum bands, efficient spectrum sharing, spectrum monitoring, and development of novel applications that leverage advanced wireless communication networks. In partnership with the private sector and other federal agencies, CISE will accelerate research in areas with potential significant impact on emerging Next-Generation (NextG) wireless and mobile communications, networking, sensing, and computing systems, with a focus on greatly improving the resiliency and intelligence of such networked systems, through the Resilient & Intelligent NextG Systems (RINGS) program and other related investments. CISE investments in at-scale research testing platforms through the Platforms for Advanced Wireless Research program will also expand engagement in important advanced wireless areas, including affordable rural broadband and autonomous aerial vehicles. These investments enable broad researcher access to large-scale research resources and accelerate the translation of innovative research outcomes in academic and government labs to societal benefits and to successful commercial products and services.

- AI: CISE, together with other NSF directorates/offices, other federal agencies, and the private sector, will increase support for AI research and development. A key focal point will be support for the National AI Research Institutes. These center-scale projects advance foundational research; conduct use-inspired research; build the next generation of talent, including at emerging research institutions; mobilize multidisciplinary groups of scientists, engineers, and educators; comprise multiple organizations working together to create significant new research capabilities; and serve as a nexus point for multisector collaborative efforts. The National AI Research Institutes will fill a major gap in America's AI research and education portfolio by accelerating AI innovations, training AI researchers and innovators, and transitioning outcomes across a range of sectors. CISE investments in AI align with the *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*.[3]
- Climate: Clean Energy Technology: CISE will support research and education projects on all sustainability topics in which advances in computing and information management are indispensable, including the areas of advanced sensing techniques; large-scale data management and analytics; optimization, modeling, simulation, prediction, and inference; intelligent systems and decision making; infrastructure design, control, and management; and human-computer interaction and social computing. Additionally, the widespread, intensive use of computing technologies introduces further sustainability challenges and motivates new approaches across the lifecycle of technology design, use, and decommission.
- Climate: USGCRP: Through its ITR division, CISE will invest funds in the NDC for Climate to federate access to advanced compute, data, software, and networking resources from multiple sources, including NSF-funded advanced computing resources, edge resources located at NSF major facilities, and capabilities deployed at other compute- and data-intensive NSF research facilities, as well as commercial cloud computing resources. The NDC for Climate will incorporate systems to curate, federate, and provide access to data from multiple sources. These approaches will advance our understanding of the Earth's climate by supporting the broad examination and reexamination of collected data, and by supporting scientific analysis of combinations of data from different sources, be they NSF-funded large facilities, resources provided by other organizations, or the data contributions of individual researchers. The NDC for Climate will further NSF's commitment to equity by democratizing access to research resources, along with the necessary support services, including outreach, on-ramping, and access to the resources made available through the NDC for Climate.
- CSGrad4US: CISE will select, recognize, and financially support early-career individuals with the demonstrated potential to be high-achieving CISE researchers and innovators, with the goal of developing the national workforce necessary to ensure the Nation's continued leadership in advancing CISE research and innovation. Through this investment, CISE aims to increase the number and diversity of domestic graduate students pursuing graduate degrees and research and innovation careers in the CISE fields—computer science, computer engineering, and/or information science—and broaden participation among groups underrepresented in these areas, including women, African Americans, Hispanics, American Indians, Alaska Natives, Native Hawaiians, Native Pacific Islanders, and persons with disabilities.
- Microelectronics and Semiconductors: CISE will support research to address fundamental science and engineering questions about the concepts, materials, devices, circuits, and platforms necessary to sustain progress in microelectronics and semiconductor technologies. Such progress is critical for emerging technologies such as AI and quantum computing and will in turn contribute

[3] www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf

to advances across all sectors of the economy, including energy, transportation, health care, and advanced manufacturing. Investments in microelectronics and semiconductor research will enable whole-of-government access to trusted and assured systems for future storage and computing paradigms.

- QIS: CISE will continue to advance quantum computing, quantum communication, and other quantum-based approaches for processing, communicating, and using information. CISE investments will specifically support novel quantum algorithms, programming languages, architectures, and circuits; simulation of quantum algorithms and systems; and designing, programming, optimizing, and testing quantum computers and systems, including through cloud-based services. Part of CISE's investments in QIS will be continuing to grow capacity within academic computer and information science departments, including cross-disciplinary and multi-department collaborations, to support advances in quantum computing and/or communication over the long term.
- SaTC: CISE will continue to lead SaTC in partnership with EDU, ENG, MPS, and SBE, investing in current and emerging areas of importance for security and privacy. These areas include the application of AI to security, security and resilience of AI systems, security implications of quantum computation and communication, information integrity, and critical infrastructure security. CISE SaTC investments will also nurture the next generation of American cybersecurity and privacy researchers and practitioners. CISE will fund programs that strengthen the national cybersecurity workforce pipeline through education, K-12 programs, and funding to universities and colleges. This funding is intended to further expand and initiate cybersecurity education programs, which improve education delivery methods for K-12 students, teachers, counselors, and post-secondary institutions and encourage students to pursue cybersecurity careers.

CISE Funding for Centers Programs

CISE Funding for Centers Programs

(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Change over FY 2021 Actual Percent
Artificial Intelligence Research Institutes (Multiple)[1]	$40.55	-	$30.50	-$10.05	-24.8%
STC: Center for Brains, Minds & Machines: The Science & the Tech. of Intelligence (CCF, IIS, ITR)	4.15	-	-	-4.15	-100.0%
STC: Center for Learning the Earth with Artificial Intelligence and Physics (IIS)	1.50	-	-	-1.50	-100.0%
Total	**$46.20**	**-**	**$30.50**	**-$15.70**	**-34.0%**

[1] In FY 2021, CISE increased investments in AI Institutes by $15.5 million over the $25.5 million originally planned. The FY 2023 Request reflects an increase of $5.0 million over the FY 2021 Current Plan.

For detailed information on individual centers programs, please see the Cross Theme Topics section of the NSF-Wide Investments chapter.

Funding Profile

CISE Funding Profile

	FY 2021 Actual Estimate	FY 2022 (TBD)	FY 2023 Estimate
Statistics for Competitive Awards:			
Number of Proposals	7,247	-	8,500
Number of New Awards	1,739	-	2,300
Regular Appropriation	1,694		2,300
ARP	45		
Funding Rate	24%	-	27%
Statistics for Research Grants:			
Number of Research Grant Proposals	7,054	-	8,300
Number of Research Grants	1,625	-	2,150
Regular Appropriation	1,582		2,150
ARP	43		
Funding Rate	23%	-	26%
Median Annualized Award Size	$166,549	-	$166,000
Average Annualized Award Size	$224,030	-	$225,000
Average Award Duration, in years	3.1	-	3.0

In FY 2023, the number of research grant proposals is expected to increase as compared to the FY 2021 Actual Estimate, and correspondingly the number of research grant awards is anticipated to increase to 2,150. The funding rate for research grants is expected to be 27 percent in FY 2023, an increase over the FY 2021 Actual Estimate. Average annualized award size and average award duration are expected to remain about the same between the FY 2021 Actual Estimate and the FY 2023 Estimate.

People Involved in CISE Activities

Number of People Involved in CISE Activities

	FY 2021 Actual Estimate	FY 2021 ARP Actual Estimate	FY 2022 (TBD)	FY 2023 Estimate
Senior Researchers	8,110	235	-	9,300
Other Professionals	1,438	63	-	1,600
Postdoctoral Associates	522	72	-	600
Graduate Students	6,150	69	-	7,000
Undergraduate Students	2,933	50	-	3,400
Total Number of People	**19,153**	**489**	**-**	**21,900**

OFFICE OF ADVANCED CYBERINFRASTRUCTURE (OAC) $252,250,000
+$21,810,000 / 9.5%

OAC Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Percent
Total	**$230.44**	**-**	**$252.25**	**$21.81**	**9.5%**
Research	**97.29**	**-**	**97.65**	**0.36**	**0.4%**
Centers Funding (total)	12.00	-	4.00	-8.00	-66.7%
Artificial Intelligence Research Institutes	12.00	-	4.00	-8.00	-66.7%
Education	**9.18**	**-**	**10.35**	**1.17**	**12.7%**
Infrastructure	**123.97**	**-**	**144.25**	**20.28**	**16.4%**
Networking and Computational Resources	121.48	-	144.25	22.77	18.7%
Mid-scale Research Infrastructure	2.49	-	-	-2.49	-100.0%

About OAC

OAC supports conceptualization, design, and implementation of the advanced research cyberinfrastructure (CI) ecosystem that is critical to advances in all areas of science and engineering research and education in the 21st century, including supporting the national response to the COVID-19 pandemic, and enabling innovations in AI, QIS, and advanced wireless, which are critical to the Nation's economy and future jobs. OAC investments also further understanding of climate science and clean-energy technologies by enabling data science, artificial intelligence and machine learning, and predictive and high-end computational modeling and simulation. Given its role across all of science and engineering, OAC works in partnership with all NSF directorates and offices as well as other CISE divisions to provide support to academic institutions, encouraging a rich and vibrant ecosystem that blends translational computer science, computational research, and research-specific CI with innovations from the private sector. Specifically, OAC investments include acquisition, integration, coordination, and operations associated with shared data, secure networking, advanced computation, scientific software and data services, and the design and development of computational and data-enabled science and engineering tools. OAC also nurtures the computational and data skills and expertise needed for next-generation science and engineering research. OAC enables researchers to address complex and multidisciplinary discovery, prediction, and innovation challenges by providing access to CI resources and services, along with secure connectivity to major national and international facilities and scientific instruments. OAC promotes innovative, robust, secure, and interoperable CI, as well as sharing and collaboration among academic research infrastructure groups, other federal agencies, international research funders, and the private sector.

OAC will continue to co-chair on behalf of NSF, the National Science and Technology Council's (NSTC) Subcommittee on the Future Advanced Computing Ecosystem (FACE). The FACE Subcommittee has developed a strategic plan describing priority areas that spur research advances in new, advanced computing architectures, systems, and services to address 21st-century scientific and technological challenges and opportunities; develop and broaden the Nation's advanced computing ecosystem including software, data, and expertise; and forge and expand partnerships. The FACE Subcommittee has also developed the National Strategic Computing Reserve (NSCR), a vision for sustaining the highly

successful COVID-19 High-Performance Computing Consortium in the longer term.

OAC also co-chairs the National Artificial Intelligence Research Resource Task Force, which is charged by Congress with developing a roadmap and implementation plan for a shared computing and data infrastructure that would provide a diverse set of researchers and students across the broad spectrum of AI research and development with access to a holistic ecosystem of resources to fuel AI discovery and innovation.

In general, about 38 percent of the OAC portfolio is available to support new grants. The remaining 62 percent supports grants made in prior years.

DIVISION OF COMPUTING AND COMMUNICATION FOUNDATIONS (CCF)

$218,570,000
+$17,620,000 / 8.8%

CCF Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Change over FY 2021 Actual Percent
Total	**$200.95**	**-**	**$218.57**	**$17.62**	**8.8%**
Research	**188.75**	**-**	**203.17**	**14.42**	**7.6%**
Centers Funding (total)	4.99	-	3.00	-1.99	-39.9%
Artificial Intelligence Research Institutes	2.50	-	3.00	0.50	20.0%
STC: Center for Brains, Minds and Machines: The Science and the Technology of Intelligence (CCF, IIS, ITR)	2.49	-	-	-2.49	-100.0%
Education	**10.35**	**-**	**13.80**	**3.45**	**33.3%**
Infrastructure	**1.85**	**-**	**1.60**	**-0.25**	**-13.5%**
National Nanotechnology Coordinated Infrastructure (NNCI)	0.60	-	0.60	-	-
Research Resources	1.25	-	1.00	-0.25	-20.0%

About CCF

CCF supports research and education activities involving the mathematical, scientific, and technological foundations of computing, communication, and information. CCF's investments enable advances in the design and analysis of algorithms, computational complexity, and mathematical modeling of systems, with attention to the efficiency, fairness, correctness, and robustness of systems including AI systems. CCF also invests in foundational research on the theoretical underpinnings of information acquisition, transmission, and processing in communication and information networks, such as sensor, advanced wireless, multimedia, and biological networks. In addition, CCF provides support for advancing the design, validation, verification and evaluation of computing hardware and software through new theories, programming languages, testing approaches, and formal methods for improving system performance, correctness, usability, reliability, and scalability. CCF investments also explore the potential impact of emerging technologies, including quantum devices and systems, neuromorphic architectures, biocomputing, synthetic biology, and nanotechnology, on various facets of computation, communication, and information that are of relevance to key priorities such as climate change and the economy.

In general, about 68 percent of the CCF portfolio is available to support new grants. The remaining 32 percent supports grants made in prior years.

DIVISION OF COMPUTER AND NETWORK SYSTEMS (CNS) $266,060,000
+$28,040,000 / 11.8%

CNS Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Percent
Total	**$238.02**	**-**	**$266.06**	**$28.04**	**11.8%**
Research	**192.25**	**-**	**226.96**	**34.71**	**18.1%**
Centers Funding (total)	9.00	-	3.50	-5.50	-61.1%
Artificial Intelligence Research Institutes	9.00	-	3.50	-5.50	-61.1%
Education	**13.36**	**-**	**16.10**	**2.74**	**20.5%**
Infrastructure	**32.41**	**-**	**23.00**	**-9.41**	**-29.0%**
Research Resources	31.41	-	23.00	-8.41	-26.8%
Mid-scale Research Infrastructure	1.00	-	-	-1.00	-100.0%

About CNS

CNS supports research and education activities that advance understanding of the fundamental properties of computer systems and networks. CNS investments produce new insights into the dynamics of complex hardware and software systems and explore new architectures for future-generation computing and communication infrastructures and services, thereby lowering barriers to innovation and enhancing economic competitiveness. These investments enable future AI, quantum computing and communication, and advanced wireless systems, as well as innovations in climate and clean energy technology. CNS-enabled systems include, but are not limited to, cyber-physical, embedded, distributed, centralized, virtualized, cloud, wireless, mobile systems, and information integrity. CNS also supports research and education activities in cybersecurity, including post-quantum cryptography, to ensure that society's ubiquitous computing and communication infrastructures deliver the quality of service they are designed to achieve without disruption, while enabling and preserving privacy, security, and trust. CNS also plays a leadership role in coordinating CISE investments in systems research infrastructure and in the development of the computing workforce of the future. CNS has fostered and continues to expand partnerships with many high-tech companies and government funding agencies to enhance support for research and education programs. CNS also supports research-based pilot projects that have the potential for scalable, sustainable, and transferable impact on communities, from small to large, and rural to urban – across the US.

In general, about 73 percent of the CNS portfolio is available to support new grants. The remaining 27 percent supports grants made in prior years.

DIVISION OF INFORMATION AND INTELLIGENT SYSTEMS (IIS)

$248,160,000
+$30,380,000 / 14.0%

IIS Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Percent
Total	**$217.78**	**-**	**$248.16**	**$30.38**	**14.0%**
Research	**202.00**	**-**	**231.96**	**29.96**	**14.8%**
Centers Funding (total)	11.38	-	10.00	-1.38	-12.1%
Artificial Intelligence Research Institutes	9.05	-	10.00	0.95	10.5%
STC: Center for Brains, Minds and Machines: The Science and the Technology of Intelligence (CCF, IIS, ITR)	0.83	-	-	-0.83	-100.0%
STC: Center for Learning the Earth with Artificial Intelligence and Physics	1.50	-	-	-1.50	-100.0%
Education	**10.71**	**-**	**14.20**	**3.49**	**32.5%**
Infrastructure	**5.06**	**-**	**2.00**	**-3.06**	**-60.5%**
Research Resources	2.06	-	2.00	-0.06	-3.0%
Mid-scale Research Infrastructure	3.00	-	-	-3.00	-100.0%

About IIS

IIS supports research and education activities that advance our knowledge of AI, data science, and human-computer interaction. The range of research topics within these areas is broad: AI includes work on knowledge representation and reasoning, machine learning, human language technologies, and computer vision; data science includes data collection and management, data integration, data mining and analytics, and informatics; and human-computer interaction includes useability, interfaces, assistive technology, and the social impacts of computing. The work supported by IIS lays the foundations for building more intelligent, human-compatible computing systems capable of advancing all sectors of the economy and society. IIS partners with other divisions, directorates, and agencies to advance diverse areas of foundational AI, data science, and human-computer interaction research across almost all areas of science, engineering, and society, including climate change and racial equity.

In general, about 70 percent of the IIS portfolio is available to support new grants. The remaining 30 percent supports grants made in prior years.

DIVISION OF INFORMATION TECHNOLOGY RESEARCH (ITR) $165,740,000 +$45,800,000 / 38.2%

ITR Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request	Change over FY 2021 Actual Amount	Change over FY 2021 Actual Percent
Total	**$119.94**	**-**	**$165.74**	**$45.80**	**38.2%**
Research	**90.87**	**-**	**150.95**	**60.08**	**66.1%**
Centers Funding (total)	8.83	-	10.00	1.17	13.3%
Artificial Intelligence Research Institutes	8.00	-	10.00	2.00	25.0%
STC: Center for Brains, Minds and Machines: The Science and the Technology of Intelligence (CCF, IIS, ITR)	0.83	-	-	-0.83	-100.0%
Education	**6.86**	**-**	**3.55**	**-3.31**	**-48.2%**
Infrastructure	**22.22**	**-**	**11.24**	**-10.98**	**-49.4%**
Research Resources	17.22	-	11.24	-5.98	-34.7%
Mid-scale Research Infrastructure	5.00	-	-	-5.00	-100.0%

About ITR

ITR provides support for transformative explorations in computer and information science and engineering research, infrastructure, and education, which are foundational for a wide range of emerging industries. These investments support emerging and high-priority areas that cut across traditional disciplinary boundaries and promise to accelerate discovery at the frontiers of the field. This includes support for foundational research on AI, QIS, particularly quantum computation and communication, and advanced wireless as well as the development of world-class research infrastructure. ITR further catalyzes research through innovative partnerships and collaborations between academia and industry. ITR will also support the development of a National Discovery Cloud (NDC) for climate that will offer large-scale democratized and equitable access to advanced compute, data, and software resources. ITR investments, often in partnership with all CISE divisions as well as NSF directorates, agencies, and industry, further research that address climate, clean energy, and equity, as well as grow our economy and jobs.

ITR, in partnership with all of the NSF directorates and research offices, will advance the HDR Big Idea by investing funds to support convergent activities that transcend the traditional disciplinary boundaries of individual NSF directorates and offices. While budget management and reporting for this investment will be the responsibility of CISE, the convergent activities will be overseen and managed collaboratively by the multi-directorate/office HDR leadership team.

In general, about 38 percent of the ITR portfolio is available to support new grants. The remaining 62 percent supports grants made in prior years.

APPENDIX A – ADVANCED COMPUTING SYSTEMS AND SERVICES PORTFOLIO

Advanced Computing Systems and Services Funding
(Dollars in Millions)

	FY 2021 Actual	FY 2022 (TBD)	FY 2023 Request
Leadership Class Computing	$19.50	-	$0.50
Advanced/Innovative Computing Systems and Services	68.56	-	82.77
Coordination and Support Services	1.52	-	14.80
Total	**$89.58**	**-**	**$98.07**

Advanced Computing Systems and Services Overview

For nearly four decades, NSF has been a recognized leader in enabling the innovative use and broad availability of a cohesive, powerful, and advanced computing ecosystem to accelerate fundamental science and engineering (S&E) research. Going forward, NSF aims to sustain America's leadership in the research, development, and broad deployment of existing as well as new advanced computing technologies, services, and skills, in part through its co-leadership of the all-of-government National Science and Technology Council (NSTC) Future Advanced Computing Ecosystem (FACE) Subcommittee efforts. Within the broad goals set for the FACE[4,5] and as further elaborated by the NSTC FACE Subcommittee, key NSF foci include fundamental and translational research to support future generations of the advanced computing ecosystem; research cyberinfrastructure (CI) including software and data services to promote cohesive platforms and interoperability for large-scale data analytics as well as modeling and simulation applications across all of S&E; and the CI expertise necessary for advancing the frontiers of CI as well as enabling S&E discovery and innovation using CI. These foci include an emphasis on a holistic approach to America's computational and data infrastructure for S&E research, spanning both human and technical dimensions, and involve forging and expanding partnerships that ensure American leadership in science, technology, and innovation. For example, during the novel coronavirus disease 2019 (COVID-19) pandemic, NSF's suite of complementary advanced computing systems and coordination services were mobilized as key contributors to the COVID-19 High-Performance Computing (HPC) Consortium, a public-private partnership that NSF helped co-found to support cutting-edge scientific research in epidemiology, virology, and microbiology, among other topics.[6]

The overall NSF advanced computing strategy and program portfolio receives guidance and input from the Advisory Committee on Cyberinfrastructure (ACCI); the Assistant Directors (AD) Council, which includes ADs and office heads from the NSF research and education directorates and offices; the Cyberinfrastructure Strategy Group, which includes senior leadership from the NSF research and education directorates and offices, and directly from the research community through multiple sources including principal investigator meetings, workshops, sessions at professional conferences,[7] community blue-ribbon studies, and Requests for Information (RFIs). In 2019, NSF funded a conference focused on the *National Cyberinfrastructure Coordination Service Conference*, which

[4] www.nitrd.gov/news/2020/Future-Advanced-Computing-Ecosystem-Strategic-Plan-Nov-2020.aspx
[5] www.nsf.gov/cise/nsci/
[6] covid19-hpc-consortium.org/
[7] See, for example, https://sc20.supercomputing.org/proceedings/bof/bof_pages/bof143.html

examined the configuration of services intrinsic to a national CI.[8] Later in the year, NSF issued a RFI asking for input on "specific data-intensive S&E research questions and challenges and the essential data-related CI services and capabilities needed to publish, discover, transport, manage and process data in secure, performant and scalable ways to enable data-intensive research."[9] Although focused primarily on data and software CI, the responses to this RFI[10] have implications for the architectures of future advanced computing systems and the services associated with maintaining and operating them. In August and September 2020, NSF sponsored the CI Workforce Development Workshop[11] focused on issues related to building and enhancing the cyberinfrastructure professional workforce. Additionally, international activities to accelerate investments in leadership-class computing, particularly in Europe and Asia, are providing additional urgency and importance for this investment strategy to ensure that the U.S. maintains its global leadership role in S&E.

In response to rapid advances in technology, changes in the capabilities and services offered by commercial interests (e.g., cloud services), and the rapid evolution of S&E research requirements, in FY 2019, NSF released a forward-looking computational ecosystem blueprint[12] and invested in three broad and complementary advanced computing areas in order to meet continually evolving needs in an agile yet predictable way. These investment areas complement each other as well as discipline-specific investments by NSF's directorates, mission-specific investments by other agencies, and cumulatively extensive, but individually smaller, investments by academic institutions at the regional and campus levels. Specifically, these areas are:

- Leadership-Class Computing, which aims to provide unique services and resources to advance the largest and most computationally intensive S&E research frontiers not otherwise possible;
- Advanced/Innovative Computing Systems and Services, which aims to provide a technically diverse, connected, and potentially future-looking advanced computing portfolio, reflecting the growing and changing use of computation and data in both the research and education processes, and capable of supporting hundreds to thousands of investigators conducting cutting-edge S&E research; and
- Coordination and Support Services, which aims to coordinate the provisioning, allocation, and operations of NSF's advanced computing resources, providing advanced assistance to the user community, supporting aggregation and federation capabilities, enabling the translation of CI research advances, and broadening participation.

In FY 2023, NSF-funded advanced computing systems and services will support the full breadth of NSF-funded S&E, including research furthering our understanding of climate science and clean-energy technologies, notably (i) data-driven approaches to assimilate heterogeneous data sets about climatology; (ii) large-scale modeling of Earth systems; and (iii) high-end simulations of renewable and alternative energy approaches, and novel materials supporting energy efficiency and sustainability.

[8] www.rti.org/publication/national-cyberinfrastructure-coordination-service-conference
[9] www.nsf.gov/pubs/2020/nsf20015/nsf20015.jsp
[10] www.nsf.gov/cise/oac/datacirfi/rfi_responses.jsp
[11] www.rcac.purdue.edu/ciworkforce2020
[12] www.nsf.gov/cise/oac/vision/blueprint-2019/nsf-aci-blueprint-v10-508.pdf

Leadership-Class Computing

Description

Leadership-class computing systems have represented a key component of NSF's computational portfolio for decades. NSF's current leadership-class computing system is Frontera, which is deployed at the Texas Advanced Computing Center (TACC) at the University of Texas at Austin (UT Austin). Frontera is one of the most powerful supercomputers in the world and is the most powerful supercomputer ever deployed on an U.S. academic campus. The system began accepting early S&E research users in May 2019 and became fully operational in October 2019. Frontera is expected to allow researchers to tackle much larger and more complex S&E applications than ever before, within and across disciplines as diverse as biology, astronomy, engineering, materials science, and the geosciences. The Frontera system offers the highest scale, throughput, and data analysis capabilities ever deployed on a U.S. university campus. In addition, Frontera's graphics processing unit (GPU) accelerates discoveries in important research areas such as deep learning and molecular dynamics.

Current Status

At its July 2018 meeting, the NSB authorized the Director to make an award to TACC for the acquisition of the Frontera system in an amount not to exceed $60.0 million over a period of five years, the first acquisition in a two-phased process. The NSB, at its May 2019 meeting, authorized the Director to make an award to TACC for the operations and maintenance (O&M) of Frontera in an amount not to exceed $60 million over a period of five years. Frontera has been in operation since September 2019 and is being actively used by the S&E research and education community across NSF and other agencies.

The July 2018 NSB resolution also authorized, pending appropriate approval associated with MREFC policies, supplemental funding to advance the design of a Phase 2 leadership-class computing facility (LCCF). In July 2019, TACC started the design and planning process for the LCCF. As noted in solicitation NSF 17-558[13] and as reported to Congress in response to the recommendations set forth in *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020*, the LCCF planning will lead to the design of a major new facility that will host a new system with a ten-fold or more time-to-solution performance improvement over the Frontera system. The Frontera system is providing S&E evaluation to inform the design of the future facility. LCCF planning will be managed and overseen according to the NSF MREFC process. The project is therefore subject to MREFC policies regarding entry and approval into the required design stages as laid out in the NSF Research Infrastructure Guide.[14] LCCF planning will continue in FY 2023 with the start of construction of the future facility anticipated in FY 2024, pending successful reviews and approvals pursuant to the NSF MREFC process.

S&E Research and Education Activities Enabled by Leadership-Class Computing

Leadership-class computing systems enable investigators across the Nation to conduct innovative research that is not otherwise possible due to demanding computing requirements. In FY 2020, NSF issued a Dear Colleague Letter[15] describing a new innovative pilot mechanism for the Nation's researchers to request access to Frontera to enable scientific and engineering research that would

[13] www.nsf.gov/pubs/2017/nsf17558/nsf17558.htm
[14] www.nsf.gov/pubs/2021/nsf21107/nsf21107.pdf
[15] www.nsf.gov/pubs/2020/nsf20018/nsf20018.jsp

not otherwise be possible without access to a leadership-class computing resource. To date, this effort has resulted in over 100 allocation awards to research teams across the country. Examples of research that were enabled by the Frontera allocation awards include the full-scale modeling of the entire hippocampus in the brain to understand neurological disorders; simulations of supermassive black hole mergers to enable future gravitational wave detection; some of the largest simulation in the world to understand the physics and conditions that cause the formation of severe tornados; and high-resolution seismic hazard modeling to improve the health and safety of the Nation's earthquake prone regions.

In addition, Frontera continues to provide important compute cycles in the all-of-nation effort in response to the COVID-19 pandemic, including as a key contributor to the COVID-19 HPC Consortium. For example, the system provided significant computing capabilities to researchers seeking to understand the fundamental infection vectors through large-scale, all-atom simulations of the SARS-CoV-2 virus, as well as tracking the epidemiology of the virus to devise better intervention strategies for preventing disease spread.

NSF-funded leadership-class computing education and outreach activities consist of projects targeting students at pre-college, undergraduate, graduate, and post-graduate levels; workshops, conferences, summer schools, and seminars; as well as industry partnership activities. These activities have enabled more than 200 education, outreach, and training projects at over 160 institutions, including institutions in the Established Program to Stimulate Competitive Research (EPSCoR) jurisdictions. An example of one of these activities is the Frontera Computational Science Fellowship program,[16] which provides a year-long opportunity for talented graduate students to compute on Frontera and collaborate with experts at TACC; this program awarded four fellowships in FY 2021.

Management and Oversight

The Frontera project is overseen by OAC's program directors and BFA's Division of Grants and Agreements staff, who receive strategic advice from the AD Council. Advice from the NSF Office of General Counsel is also sought, as necessary. Planning for the LCCF system is coordinated with the Large Facilities Office and the Division of Acquisition and Cooperative Support in BFA and will be reviewed in accordance with NSF's major facilities policies and procedures. The NSB receives updates on any major changes in risk assessments, which are reviewed annually by an external panel. Risks monitored during the operational phase of a project include system security, performance, reliability, usability, project management, and other factors that could reduce the overall scientific impact.

Advanced/Innovative Computing Systems and Services

Description

NSF funds the acquisition and operation of nationally available Advanced/Innovative Computing Systems and Services that, in aggregate, are forward-looking, connected, and technically diverse, and reflect changing and growing use of data-intensive computation in both the research and education processes. At the same time, they are intended to enable discoveries at a computational scale beyond the reach of an individual or regional academic institution.

Deployed systems currently serve as a cohesive set of allocable resources within the eXtreme Digital

[16] frontera-portal.tacc.utexas.edu/fellowship/

(XD) integrated services infrastructure, which is described in the following section. Awards are generally made as two parts: an acquisition and deployment award, which may be the result of a competitive or a renewal proposal; and a separate award for O&M following deployment. When an award is made, the awardee institution issues subawards to vendors and/or other organizations for acquisitions and services, as necessary. Expenditures are contingent on successful completion of deployment milestones. These systems are also accessible via the Partnership to Advance Throughput Computing (PATh) project.[17]

Current Status

In FY 2016, NSF awarded *Stampede 2: The Next Generation of Petascale Computing for Science and Engineering* to TACC, enabling the acquisition and deployment of Stampede 2. Stampede 2 serves as the primary national resource for approximately 7,000 academic researchers, complements other national advanced computing systems and services, and provides capabilities beyond the reach of individual campuses and regional resources. Stampede 2 was fully deployed as a production resource by the end of 2018 and is expected to continue operations through June 2023. This includes a technical upgrade awarded in FY 2021 to extend operations, partially upgrade the processor architecture, and explore pilot high throughput computing allocations via the PATh project.[18]

In addition, beginning in FY 2019, NSF made a series of investments in advanced/innovative computing systems and services to foster an integrated CI ecosystem that addresses the growing scale and diversity of the S&E community, the changing nature of S&E research requirements, and the rapidly evolving technology and services landscape, with the overarching goal of supporting the full range of computational- and data-intensive research across all S&E domains. Specifically, NSF issued the *Advanced Computing Systems and Services (ACSS): Adapting to the Rapid Evolution of Science and Engineering Research* solicitation[19] in FY 2019, with the first cohort of three awards running from FY 2019 to FY 2024,[20] followed by a second cohort of five awards running from FY 2020 to FY 2025,[21] and a third cohort of 2 awards running from FY 2021 to FY 2026.[22]

The ACSS solicitation called for investments in two categories:

- Category I, Capacity Systems: production computational resources maximizing the capacity provided to support the broad range of computation and data analytics needs in S&E research; and
- Category II, Innovative Prototypes/Testbeds: innovative forward-looking capabilities deploying novel technologies, architectures, usage modes, etc., and exploring new target applications, methods, and paradigms for S&E discoveries.

In the FY 2019 ACSS competition, two Category I awards were made to the Pittsburgh Supercomputing Center (PSC), and the San Diego Supercomputer Center (SDSC) at UCSD; and one Category II award was made to the State University of New York (SUNY) at Stony Brook. In the FY 2020 ACSS competition, three Category I awards were made to Indiana University, Purdue University, and University of Illinois at Urbana-Champaign (UIUC); and two Category II awards were made to SDSC and PSC. In the FY 2021

[17] www.nsf.gov/awardsearch/showAward?AWD_ID=2030508
[18] www.nsf.gov/awardsearch/showAward?AWD_ID=2030508
[19] nsf.gov/funding/pgm_summ.jsp?pims_id=503148
[20] www.nsf.gov/pubs/2019/nsf19534/nsf19534.htm
[21] www.nsf.gov/pubs/2019/nsf19587/nsf19587.htm
[22] www.nsf.gov/pubs/2020/nsf20606/nsf20606.htm

ACSS competition, two Category II awards were made to the San Diego Supercomputer Center (SDSC) at UCSD and Texas A&M University. Given interruptions to supply chains resulting from the COVID-19 pandemic, some scheduled deployments were slightly delayed, but disruptions to the relevant S&E research communities have been minimal. When fully deployed, the suite of Category I systems will include the following:

- *Expanse*: Located at SDSC this system will be operational from FY 2021 through FY 2024. Expanse is a large-capacity, data-focused system supporting increasingly diverse, complex, and expanding research across multiple S&E disciplines within the "long tail" of science.
- *Bridges 2*: Located at PSC this system will be operational from FY 2021 through FY 2024. Bridges 2 integrates AI-based analytics capabilities with the technical capacity to execute data- and computationally-intensive research in broad, cross-cutting manners, enabling advances across a range of S&E research and education.
- *Anvil*: Located at Purdue University, a new service provider within the NSF ecosystem of advanced computing systems, Anvil will be operational from FY 2022 through FY 2025. Anvil will be a composable system with an expansive portfolio of S&E-focused interfaces, programming environments, and advanced capabilities to support research and education.
- *Delta*: Located at the UIUC, Delta is expected to be operational from FY 2022 through FY 2025. Delta will be a large-capacity, balanced computational resource supporting traditional computational methods combined with rapidly evolving and expanding AI-based techniques and advanced data science methods to advance S&E research and education. Transition to production operations for Delta following a successful acceptance review is expected in 2022.
- *Jetstream 2*: Located at Indiana University Jetstream 2 is expected to be operational from FY 2022 through FY 2025. Jetstream 2 will provide a nationally distributed, large-capacity, cloud-enabled computational resource supporting diverse S&E-focused "on-demand" access modes and utilization models to be available across research and education. Jetstream 2 is expected to transition to production operations following a successful acceptance review in 2022.

In addition, the Category II, or Testbed-Prototype Systems, comprise:

- *Ookami*: Located at SUNY at Stony Brook, this prototype will be operational through FY 2024. NSF will evaluate the utility of the system and determine whether it can be integrated into the suite of production services. Ookami incorporates processors originally developed to lead Japanese national efforts[23] towards future computing to advance U.S.-based S&E research and education.
- *Neocortex*: Located at PSC, this prototype will be operational through May 2025. NSF will evaluate the utility of the system and determine whether it can be integrated into the suite of production services. Neocortex will deploy a novel AI-focused processor architecture in a high-performing system design supporting very high-scale, complex analytics challenges across S&E research and education.
- *Voyager*: Located at SDSC, this prototype will be operational through May 2025. NSF will evaluate the utility of the system and determine whether it can be integrated into the suite of production services. Voyager will integrate AI/ML/deep learning-focused components to advance S&E research and education.
- *National Research Platform (NRP):* Located at SDSC, with partners at University of Nebraska, Lincoln (UNL) and the Massachusetts Green High Performance Computing Center (MGHPCC), this prototype will be operational through May 2026. The prototype NRP will deploy a distributed testbed architecture including a high-performance subsystem at SDSC; two Graphics Processing

[23] www.r-ccs.riken.jp/en/fugaku/project

Unity (GPU) subsystems at UNL and MGHPCC; and low latency high bandwidth research and education networking.

- *Accelerating Computing for Emerging Sciences (ACES)*: Located at Texas A&M University, this prototype system will be operational through September 2026. ACES will deploy a novel composable system architecture with the flexibility to aggregate various components on an as-needed basis to solve problems previously not addressable by researchers.

NSF will evaluate the utility of the above listed Category II, or Testbed-Prototype Systems and determine whether they can be integrated into the suite of production services.

S&E Research and Education Activities Enabled by Advanced/Innovative Computing Systems and Services

The ecosystem of advanced/innovative computing systems and services is enabling new, world-leading, and transformative advances across the breadth of S&E research, in the integration of research and education, and in broadening participation in S&E by underrepresented groups. It is enabling new collaborations across public and private sectors to advance American security and economic competitiveness. These advances are made possible by providing researchers and educators with access to world-leading computational systems and services beyond what is typically available on most campuses. Providing access includes providing the expertise, interfaces, consulting support, and training necessary to facilitate use of the systems and services. This activity is central to America achieving the full potential of complementary investments by NSF, other federal agencies, and academic institutions in computing infrastructure.

Management and Oversight

OAC's program directors provide direct oversight during both the acquisition and O&M awards. Formal reporting consists of quarterly and annual reports, which are reviewed by the program directors.

Awards for advanced/innovative computing system and services are managed under cooperative agreements that include management structures, milestones, spending authorization levels, and review schedules. Each awardee is responsible for the satisfactory completion of milestones prior to NSF authorization of spending. Progress is assessed with the aid of annual external reviews. In addition, each project is required to have a project management plan.

Any activity of this nature and at this scale comes with a certain element of risk. The review process, conducted prior to award, analyzes the risks as presented in the proposal and identifies any additional risks that should be considered. During the award process, risks are identified and analyzed, and a mitigation plan is created and followed. One of the activities that are a part of the periodic NSF external reviews conducted by an external panel of experts is to revisit and reassess the risk and make recommendations as deemed necessary. In the case of projects that involve an acquisition, project risks are generally substantially reduced after deployment. Thus, the pacing of the acquisitions and deployments for such projects provides balance in the overall risk portfolio for the program.

Milestone-driven reviews occur during the acquisition award, typically with an external review prior to deployment. Annual reviews, conducted by an external panel of expert reviewers and managed by OAC program directors, are performed during the operational phase of each project.

Coordination and Support Services

Description

NSF's investments in coordination and support services, as exemplified by the XD integrated services infrastructure, add value to the NSF advanced/innovative computing systems and services by coordinating allocations and access to the systems and services, providing advanced assistance to the user community, and broadening participation. The XD program's shared services model for coherently and efficiently providing researchers with both access to and expertise for diverse, dynamic, and distributed resources is a cornerstone of the American advanced computing ecosystem; enabling the connection between individual campuses and national resources is an essential aspect.

XD enables and supports leading-edge scientific discovery and promotes science and technology education. The program encourages innovation in the design and implementation of an effective, efficient, increasingly virtualized approach to the provisioning of high-end digital services, while ensuring that the infrastructure continues to deliver high-quality access for the many researchers and educators who use it in their work.

XD shared services consist of several interrelated parts: allocation of resources to computational and data research projects; advanced user assistance; training, education, and outreach; architecture and operation of an integrated digital services infrastructure; metrics services; and overall coordination. These elements are designed and implemented in a way that is clearly tied to the requirements of the S&E research community, using a flexible methodology that permits the architecture to evolve in response to changing community needs and that presents individual users with a common environment regardless of where the resources or researchers are located.

For researchers requiring high-throughput computing – computing that can be characterized by executing large numbers of tasks over a long period of time - The Partnership to Advance Throughput Computing (PATh) makes Distributed High Throughput Computing (dHTC) capacity available to researchers through a fabric of services. These services enable the federation of resources into an effective source of computing capacity for a wide spectrum of science applications. PATh supports single-PIs and collaborative science groups across science and engineering disciplines to join the cohort of international physical science collaborations who have leveraged the dHTC paradigm for decades.

Current Status

Two awards are currently active within the XD program: the eXtreme Science and Engineering Discovery Environment (XSEDE) and the XD Metrics Service (XMS). An upcoming suite of six independent, yet highly coordinated, awards planned for FY 2022 will replace XSEDE and XMS by establishing the next generation of Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS) (described below). Continued operation of XSEDE and XMS awards during a 6-month overlap period with ACCESS is intended to provide a smooth transition to services with a new organizational structure with minimal disruption to the community.

The current XSEDE award to UIUC was renewed in September 2016, continuing the prior XSEDE award for another five-year period. This five-year award has been extended for a sixth year, based on a very successful site review. Within the current XSEDE project, there are 18 partners engaged via subawards to the University of Tennessee at Knoxville (National Institute for Computational Sciences), CMU and

University of Pittsburgh (PSC), UT Austin (TACC), UCSD (SDSC), University of Chicago, Indiana University, Purdue University, Shodor Education Foundation, Ohio Supercomputer Center, Southeastern Universities Research Association, Cornell University, National Center for Atmospheric Research, Georgia Institute of Technology, University of Georgia, Oklahoma University, University of Southern California, University of Arkansas, Notre Dame, and Internet 2.

The current XMS award was made in FY 2015 to SUNY at Buffalo. This award provides metrics services allowing measurement and monitoring of key operational data from XSEDE services and the advanced computing/innovative systems and services portfolio. The mid-project external site review of the XMS project took place in June 2018 and continued operations were authorized based on the successful outcome of that review.

The main award, "Partnership to Advance Throughput Computing (PATh)", is a five-year award to the University of Wisconsin-Madison. Within the award, six partners are engaged through sub-awards: Indiana University, Information Sciences Institute (USC), Morgridge Institute for Research, University of California San Diego, University of Chicago, and University of Nebraska-Lincoln. The award is now in its 2nd year.

NSF has outlined plans for a fabric of national CI coordination services in a blueprint document released in FY 2020.[24] This blueprint is based on findings from the NSTC FACE Subcommittee, guidance from ACCI and advisors, responses to an RFI, and feedback from engagement with the community about the structure and composition of future coordination efforts. Following the blueprint, NSF issued the *Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS)* and *ACCESS - Coordination Office (ACCESS-ACO)* solicitations[25-26] in early FY 2021. Awards for ACCESS services and the ACCESS-ACO will be made during FY22.

S&E Research and Education Activities Enabled by Coordination and Support Services

Coordination and support services, as exemplified by XD and to be expanded by ACCESS awardees, enable transformative advances in S&E research, in the integration of research and education, and in broadening the participation of underrepresented groups in S&E. These advances are accomplished by providing researchers and educators with coherent and highly usable access to digital resources beyond those typically available on most campuses, together with the interfaces, consulting, advanced user support, and training necessary to facilitate their use.

Coordinated access to advanced/innovative computing systems and services enables researchers to efficiently manipulate, analyze, visualize, and share extremely large amounts of distributed digital information from simulations, sensors, and experiments. The ACCESS awards will enable the cyberinfrastructure ecosystem resources and CI professionals to innovate and evolve in sync with S&E research and education needs and opportunities. External communication, outreach, and community-building efforts by the ACCESS awardees will broaden the participation of individuals and communities that have been underserved by the national CI ecosystem.

The fabric of coordination and support services for the advanced CI ecosystem deliver tools and

[24] www.nsf.gov/cise/oac/vision/blueprint-2019/nsf-aci-blueprint-services.pdf
[25] www.nsf.gov/pubs/2021/nsf21555/nsf21555.htm
[26] www.nsf.gov/pubs/2021/nsf21556/nsf21556.htm?org=NSF

democratized access for researchers seeking resources, such as those described above, but also enable scientific collaborations for geographically distributed teams. In doing so, these services facilitate dynamic access to digital resources and experimental testbeds within and across university campuses, as well as government laboratories. These services also support the integration research software and data with CI resources. Human-in-the-loop expert services and widely available training materials reduce barriers to the use of advanced digital systems by the research and education communities, thereby promoting enhanced productivity. For example, the XSEDE platform has provided the basis for coordination and resource allocation among the more than 40 members of the COVID-19 HPC Consortium, and the team continues to provide essential services to support end users.

Monitoring and measurement services collect multi-dimensional data on advanced CI ecosystem usage statistics, users, and the computing resources' performance. They have also deployed CI measurement and optimization tools, namely XDMoD (XD Metrics on Demand) and its open-source counterpart, Open XDMoD, which are in use worldwide for advanced CI monitoring and reporting in academia and industry. Ongoing investments in these tools will enable the exploration of novel usage modes for advanced/testbed computing systems, integration with data repositories, instrumentation, and network performance. The immediate users of these methods and tools are the providers of NSF-supported advanced computing systems and services. However, both the tools and data are publicly available and used by researchers, academic research computing center administrators, federal agencies, and industry seeking to optimize performance and forecast capacity demand.

Management and Oversight

OAC's program directors oversee the advanced CI ecosystem services and support projects. Project management is supported by guidance from an external advisory board, service provider councils, and ongoing formal and informal engagement with stakeholder communities. OAC's oversight of projects includes participation in weekly teleconferences with senior personnel of awardee teams, quarterly briefings, and regularly scheduled planning sessions such as the allocation requests review meetings. Formal reporting consists of quarterly and annual reports, which are reviewed by the program directors. Projects participate annually in site reviews or virtual reviews conducted by external panels of expert reviewers and managed by OAC program directors. Each award is managed under a cooperative agreement with tailored terms and conditions, including an approved Project Execution Plan detailing management structure, milestones, deliverables, risk management, reporting of spending levels over time, and a review schedule. Each awardee is responsible for the satisfactory completion of milestones prior to NSF authorization of spending.

The PATh award is actively managed through monthly project meetings, monthly reports, monthly updates on goals and milestones, quarterly reports, annual reports, and a Project Execution Plan that is updated at least once per year.