DIRECTORATE FOR COMPUTER AND INFORMATION SCIENCE AND ENGINEERING (CISE) $893,350,000 -$654,000 / -0.1%

CISE Funding
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Percent
Advanced Cyberinfrastructure (ACI)	$207.59	$212.29	$212.29	-	-
Computing and Communication Foundations (CCF)	178.02	185.19	185.19	-	-
Computer and Network Systems (CNS)	211.03	220.40	220.40	-	-
Information and Intelligent Systems (IIS)	176.23	185.18	185.19	0.01	0.0%
Information Technology Research (ITR)	85.25	90.95	90.29	-0.66	-0.7%
Total, CISE	**$858.13**	**$894.00**	**$893.35**	**-$0.65**	**-0.1%**

Totals may not add due to rounding.

About CISE

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; to promote understanding of the principles and uses of advanced computer, communications, and information systems in service to society; and to contribute to universal, transparent, and affordable participation in a knowledge-based society. CISE supports ambitious long-term research and research infrastructure projects within and across the many sub-fields of computing, as well as cyberinfrastructure for all areas of science and engineering; contributes to the education and training of computing professionals; and, more broadly, informs the preparation of a U.S. workforce with computing and computational competencies essential to success in an increasingly competitive global market.

CISE's FY 2015 Budget Request is shaped by the following NSF-wide priorities – Cyber-enabled Materials, Manufacturing, and Smart Systems (CEMMSS), which includes Advanced Manufacturing, Cyber-Physical Systems, and the National Robotics Initiative; Clean Energy Technology, which includes Science, Engineering, and Education for Sustainability (SEES); Cyberinfrastructure Framework for 21st Century Science, Engineering, and Education (CIF21); NSF Research Traineeship (NRT); and Secure and Trustworthy Cyberspace (SaTC). Progress in foundational research and education in these areas is vital to addressing key national challenges, spurring innovation, increasing productivity, securing critical infrastructure, improving data analysis and sharing, developing the next generation of computing and computational scientists, and, more generally, promoting economic growth.

CISE continues to provide leadership for the multi-agency Subcommittee on Networking and Information Technology Research and Development (NITRD), which is co-chaired by the CISE Assistant Director. All research, education, and research infrastructure projects supported by CISE enrich the agency's NITRD portfolio. As noted by the President's Council of Advisors on Science and Technology (PCAST) in its *Report to the President and Congress - Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology* (January 2013)[1], advances in Networking and Information Technology (NIT) are key drivers of U.S. economic competitiveness. Essentially all practical applications of Information Technology (IT) are based on ideas and concepts that emerged from

[1] www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-nitrd2013.pdf

investments in basic computing research. These fundamental ideas and concepts have enabled innovative products and applications that now permeate all areas of modern life. IT is integral to the Nation's economy, as it drives discovery and innovation in many other areas, e.g., frontiers of scientific research, advanced manufacturing, education and workforce development, health and wellness technologies, sustainability and energy science, transportation, national and homeland security research, and public and private organizational effectiveness and efficiency. Innovation in IT will remain an essential and vital force in productivity gains and economic growth in both the manufacturing and service sectors for many years to come, positioning NSF and CISE as central and essential actors in improving the Nation's economic outlook and advancing a highly trained, technologically astute workforce.

CISE provides about 87 percent of the federal funding for basic research at academic institutions in computer science.

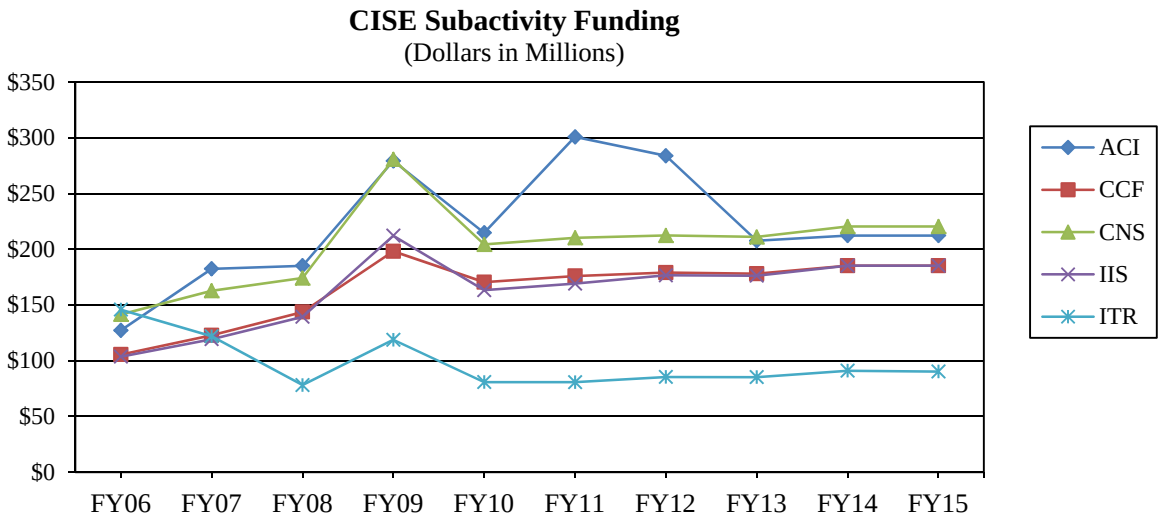

FY 2015 Summary by Division

- ACI's FY 2015 Budget Request is focused on maintaining and leveraging investments in existing programs in computational science, software, data, networking, and cybersecurity, and providing leadership in the NSF-wide CIF21 activity through programs such as Data Infrastructure Building Blocks (DIBBs), Software Infrastructure for Sustained Innovation (SI^2), EarthCube, Computational- and Data-Intensive Science & Engineering (CDS&E), and the CIF21 emphasis area in NRT. The goal of CIF21 is to accelerate and transform the progress of scientific discovery and innovation by providing cyberinfrastructure – including new data infrastructure and computational frameworks – that enable novel functionalities and capabilities in data-enabled, computational science and engineering. ACI continues to support other cross-disciplinary activities, including transitioning discoveries into practice in the SaTC program, and participating in the NSF Innovation Corps (I-Corps) and Critical Resilient Interdependent Infrastructure Systems and Processes (CRISP) programs. ACI remains responsible for providing national resources and instruments to facilitate collaborations and greater data sharing across research communities. ACI-supported infrastructure will be used to address some of the most difficult and complex research problems in all areas of science and engineering. For example, researchers can now take advantage of two new significant ACI-supported computational resources that became available in FY 2013 – Blue Waters and Stampede.

- CCF's FY 2015 Budget Request is focused on maintaining support for its core programs and NSF-wide priority areas. CCF will continue to support CIF21 through investments in Big Data foundational research, including algorithms and software tools for managing massive amounts of heterogeneous, complex data; new functional capabilities in support of highly parallel computing, and multi-core and multi-machine data management systems. CCF will increase its investments in Clean Energy Technology, with a focus on the computational methods and models necessary to attain a sustainable future. The division will continue supporting the SEES Cyber-Enabled Sustainability and Engineering (CyberSEES) program as well. CCF will continue to invest in eXploiting Parallelism and Scalability (XPS) as part of its core programs. CCF continues to support foundational research in SaTC, including new theories, models, methods, architectures, and tools that aim to achieve security-aware computing, self-healing hardware, and self-protecting software. As part of the National Nanotechnology Initiative (NNI), CCF will focus on research in nanoscale devices and systems.

- CNS's FY 2015 Budget Request is focused on maintaining support for its core programs, as well as on providing support for NSF-wide priority areas. In partnership with the other CISE divisions and the Directorates for Education and Human Resources (EHR), Engineering (ENG), Mathematical and Physical Sciences (MPS), and Social, Behavioral and Economic Sciences (SBE), CNS will continue to lead the SaTC program. CNS will maintain CISE's national leadership in developing the scientific foundations of cybersecurity as part of the Comprehensive National Cybersecurity Initiative (CNCI). Also, in partnership with other CISE divisions and NSF directorates, CNS will support CEMMSS through investment in research in Advanced Manufacturing (AM), Cyber-Physical Systems (CPS), and the National Robotics Initiative (NRI). Additionally, CNS will maintain its support for CIF21 through funding Big Data research on pervasive computing, and large-scale data management systems. With EHR, CISE will continue to support the Science, Technology, Engineering, and Mathematics, including Computing Partnerships (STEM-C Partnerships) program in FY 2015. CNS will continue its support for mid-scale network infrastructure, transitioning its investments in the Global Environment for Network Innovations (GENI) project into the NSFCloud, allowing for experimentation in future large-scale distributed computing resources not possible elsewhere.

- IIS's FY 2015 Budget Request maintains support for its core programs and NSF-wide priority areas. IIS will participate in CEMMSS through leadership of NRI, in partnership with the National Institutes of Health (NIH), National Aeronautics and Space Administration (NASA), and Department of Agriculture (USDA), as well as with other NSF directorates, including ENG, EHR, and SBE. NRI will accelerate the development and use of robots in the United States that work beside, or cooperatively with, people. IIS will participate in CIF21, and will provide leadership in Big Data analytics, including funding for new approaches to data mining, machine learning, knowledge extraction, visualization, predictive modeling, and automated discovery. IIS continues to lead a joint NSF – NIH program, Smart and Connected Health (SCH), partnering with ENG and SBE, and the other CISE divisions. IIS will co-lead with EHR the Cyberlearning and Future Learning Technologies (Cyberlearning) program. This program aims to integrate advances in technology with advances in understanding how people learn, with a focus on online learning environments. IIS also will increase its investments in cognitive science and neuroscience, building on investments in computational neuroscience and foundational research programs to advance understanding of brain functions.

- ITR's FY 2015 Budget Request supports emerging high-priority areas of potentially transformative research. Through increased investments in I-Corps, ITR will build on foundational research and guide the output of scientific discoveries in the development of technologies, products, and processes that benefit society. ITR will continue to invest in the Expeditions in Computing program, which encourages researchers to identify the compelling ideas that promise transformations in computing

and information sciences for years to come. ITR will invest in multi-disciplinary research networks, aiming to build communities across emerging areas of research and education. ITR will continue its investments in mid-scale infrastructure, extending virtualization beyond the network (i.e., in GENI) to large-scale, interconnected computing resources by investing in mid-scale prototypes for an NSF Cloud research infrastructure. Through US Ignite, ITR will continue to expand and research-enable U.S. campuses, advance networking and systems research through experimentation and explorations at scale, and jumpstart gigabit application development and deployment.

Major Investments

CISE Major Investments

(Dollars in Millions)

Area of Investment	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Change Over FY 2014 Estimate Percent
Advanced Manufacturing	$35.27	$39.63	$37.75	-$1.88	-4.7%
CAREER	50.12	44.37	44.88	0.51	1.1%
CEMMSS	64.80	85.00	81.50	-3.50	-4.1%
CIF21	57.03	85.00	80.00	-5.00	-5.9%
Clean Energy Technology	18.00	18.00	20.52	2.52	14.0%
Cognitive Science and Neuroscience	-	3.50	5.65	2.15	61.4%
I-Corps	4.60	8.00	10.00	2.00	25.0%
NSF Research Traineeship (NRT)[1]	10.25	6.89	7.59	0.70	10.2%
SEES	13.85	11.00	11.00	-	-
SaTC	59.00	70.00	67.00	-3.00	-4.3%

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

[1] The FY 2013 Actual represents Integrative Graduate Education and Research Traineeship (IGERT) program funding. Outyear commitments for IGERT are included in the NRT line and are $330,000 in FY 2014 and $900,000 in FY 2015.

- Advanced Manufacturing: As part of CEMMSS, CISE, in partnership with ENG and MPS, will invest in research that integrates ubiquitous sensors, computational tools, and highly connected cyber-physical systems in smart processing and advanced manufacturing systems. This investment will result in higher quality products with greater efficiency and sustainability produced by the factories of the future. CISE will also invest in basic research to advance robotics technology, enabling new functionalities and providing the next-generation of products and services in various industries. This includes co-robots that work alongside or cooperatively with people in manufacturing environments to increase their productivity, performance, and safety.

- CAREER: This program invests in the integration of research and education of early-career researchers and contributes to the development of future generations of computer and information scientists and engineers, as well as computational scientists across all areas of science and engineering.

- CEMMSS: CISE, in partnership with BIO, ENG, and MPS, aims to establish a scientific basis for engineered systems interdependent with the physical world and social systems, synthesize multi-disciplinary knowledge to model and simulate systems in their full complexity and dynamics, and

develop a smart systems technology framework. As part of CEMMSS, the CPS program, funded jointly with ENG and two other federal agencies —Department of Homeland Security (DHS) and Department of Transportation (DOT) — promises to accelerate advances in 21st century smart engineered systems. CEMMSS also includes CISE investments (along with EHR, ENG, and SBE) in the multi-agency NRI (with NASA, NIH and USDA). As part of this research activity, CISE will synergize investments across multiple research communities and programs in order to transform static systems, processes, and edifices into adaptive, pervasive smart systems with embedded computational intelligence that can sense, adapt, and react.

- CIF21: CISE will lead CIF21, in support of advances in the Big Data research program and through investments in the DIBBS and SI^2 programs. Big Data research will focus on core scientific and technological means of managing, analyzing, visualizing, and extracting useful information from large, distributed, and heterogeneous data sets as well as applications in specific research domains. DIBBs aims to develop, implement, and support new cyberinfrastructure to store and manage the diversity, size, and complexity of current and future data sets and streams. To advance new computational infrastructure, SI^2 promises to advance new paradigms and practices in the development and use of robust, reliable, usable, and sustainable software.

- Clean Energy Technology: CISE will support foundational research in energy-intelligent computing; the development of new theory, algorithms, and design principles to investigate energy versus computation and communication tradeoffs; and the scalability and sustainability of smart energy production software and hardware. CISE research on clean energy is partially supported via investments in SEES.

- Cognitive Science and Neuroscience: In collaboration with other directorates and offices, CISE will support new projects focused on understanding the brain. In particular, CISE will support projects that develop computational approaches for investigating neural and behavioral plasticity in response to external changes. CISE investments will enable the research needed to integrate computational models across scales; the development of innovative neurotechnologies to monitor brain function; and the expansion of the capacity of neuroscience infrastructure to integrate data across levels of analysis from molecular to behavioral scales.

- I-Corps: In FY 2015, CISE will continue to support I-Corps Teams, Sites, and Nodes to further build, utilize, and sustain a national innovation ecosystem that continues to augment the development of technologies, products, and processes that benefit the Nation. I-Corps Sites are funded at academic institutions that already have existing innovation or entrepreneurial units to enable them to nurture and support multiple, local teams to transition their ideas, devices, processes or other intellectual activities into the marketplace. The I-Corps Nodes' goal is to establish regional nodes to provide training to I-Corps Teams; develop tools and resources that will impact and expand the benefits of the entire I-Corps program within a two to three year timeframe; and identify and pursue longer-term research and development projects. CISE's investment will seek to identify NSF-funded researchers who will receive additional support in the forms of mentoring and funding to accelerate innovation that can attract subsequent third-party investment.

- NSF Research Traineeship: CISE will continue its support of the CIF21 emphasis area within the NSF Research Traineeship program, designed to encourage the development of new scalable and potentially transformative models of STEM graduate training that ensure graduate students develop the skills, knowledge, and competencies needed to pursue a range of careers within and outside academe.

- SEES: CISE will invest in CyberSEES, which aims to advance interdisciplinary research in which the science and engineering of sustainability are enabled by advances in computational- and data-intensive research and education.

- SaTC: NSF continues to align its cybersecurity investments (including investments from EHR, ENG, MPS, and SBE) with the national cybersecurity strategy, *Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program*.[2] SaTC aims to support scientific foundations, induce change, maximize research impact, and accelerate the transition of advances to practice. This investment also includes support for CNCI. CISE is collaborating with EHR to support cyber-secure workforce development to enable a growing pipeline of researchers and educators, and to develop a citizenry that understands the security and privacy of the digital systems on which society depends.

CISE Funding for Centers Programs and Facilities

CISE Funding for Centers Programs
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Percent
Total, Centers Programs	**$10.57**	**$14.69**	**$10.00**	**-$4.69**	**-31.9%**
STC: Team for Research in Ubiquitous Secure Technology (CCF)	3.32	2.66	-	-2.66	-100.0%
STC: Center for Science of Information (CCF)	5.00	5.00	5.00	-	-
STC: Center for Brains, Minds and Machines: the Science and the Technology of Intelligence	-	5.00	5.00	-	-
SLC: Pittsburgh Science of Learning Center - LearnLab (ITR)	2.25	2.03	-	-2.03	-100.0%

Totals may not add due to rounding.

For detailed information on individual centers, please see the NSF-Wide Investments chapter.

- The CISE-supported Science and Technology Center (STC): Team for Research in Ubiquitous Secure Technology (TRUST) at the University of California at Berkeley will sunset in FY 2015 (-$2.66 million to a total of zero).

- CISE will provide the fifth year of funding for the STC: Center for Science of Information at Purdue University. The goal of this center is to develop a new science of information that incorporates common features associated with data/information, such as space, time, structure, semantics, and context, that are not addressed by earlier mathematical theories, e.g., data obfuscation and hiding techniques that enhance robustness and the principles of redundancy and fault tolerance found in natural systems.

- CISE will provide the third year of funding for the STC: Center for Brains, Minds and Machines: the Science and the Technology for Intelligence. The Center has five main research themes: circuits for intelligence; the development of intelligence in children; social intelligence; the integration of visual, motor, language, and social intelligence; and theoretical aspects of intelligence.

[2] www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf

- LearnLab, formerly known as the Pittsburgh Science of Learning Center (SLC) for Robust Learning, will sunset in FY 2015 (-$2.03 million to a total of zero).

CISE Funding for Facilities

(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Change Over FY 2014 Estimate Percent
Total, Facilities	**$0.60**	**$0.60**	**$0.60**	-	-
National Nanotechnology Infrastructure Network (CCF)	0.60	0.60	0.60	-	-

Totals may not add due to rounding.

For detailed information on individual facilities, please see the Facilities chapter.

Summary and Funding Profile

CISE supports investments in core and interdisciplinary research and education, as well as in computing research infrastructure.

In FY 2015, the number of research grant proposals is expected to increase by approximately 4.9 percent compared to the FY 2014 Estimate. CISE expects to award approximately 1,430 research grants in FY 2015. Average annualized award size and average award duration are expected to remain constant between the FY 2014 Estimate and FY 2015 Estimate.

Funding for research infrastructure represents 17 percent of the CISE Request. Most of CISE's research infrastructure support is for High Performance Computing (HPC) (see Appendix A for more information on the HPC portfolio).

CISE Funding Profile

	FY 2013 Actual Estimate	FY 2014 Estimate	FY 2015 Estimate
Statistics for Competitive Awards:			
Number of Proposals	7,821	8,100	8,500
Number of New Awards	1,616	1,680	1,680
Funding Rate	21%	21%	20%
Statistics for Research Grants:			
Number of Research Grant Proposals	7,484	7,750	8,130
Number of Research Grants	1,373	1,430	1,430
Funding Rate	18%	18%	18%
Median Annualized Award Size	$161,237	$165,000	$165,000
Average Annualized Award Size	$204,242	$210,000	$210,000
Average Award Duration, in years	2.9	3.0	3.0

Program Monitoring and Evaluation

Committees of Visitors (COV)

- In early FY 2015, CISE plans to hold a Committee of Visitors (COV) review, which will examine and assess the quality of the CISE merit review process.

Science and Technology Policy Institute (STPI) Reports and Evaluations

- CISE established a contract with the Science and Technology Policy Institute (STPI) to conduct program evaluation feasibility studies for the SaTC and CEMMSS programs. These feasibility studies will examine baseline portfolio investments and identify metrics to measure progress toward program goals. They are a part of a broader effort to develop a plan for impact assessments of SaTC and CEMMSS investments. A contract was established and a kick-off meeting was held at the end of the fourth quarter of FY 2012. The preliminary work to identify baseline evaluation metrics was conducted in FY 2013, and it is anticipated that the program evaluation analyses will begin once a contract is put into place in early FY 2015. Yearly program-wide assessments will be presented to the CEMMSS and SaTC working groups and to NSF senior management.

STEM Evaluation

- Evaluation is also a vital part of CISE's STEM education programs. Each of the STEM-C Partnership projects managed by CISE has a rigorous research and/or evaluation plan designed to guide project progress and measure its impact. These plans include descriptions of the instruments and metrics that are to be used. Across the STEM-C Partnerships portfolio managed by CISE, a set of common metrics and a design for evaluation instruments are being developed. The initial design is complete, and the contractor has had the first of two face-to-face meetings with the individual project evaluators to gain their input and cooperation. The first program evaluation will be initiated under a new contract to be negotiated by early FY 2015.

Reports

- CISE funded the National Academy of Sciences (NAS) Computer Science and Telecommunications Board (CSTB) to study the IT innovation ecosystem and to assess the long-term economic impacts of CISE investments. The report, *Assessing the Impacts of Changes in the Information Technology R&D Ecosystem*,[3] published in 2009, includes an in-depth articulation of the creation of almost 20 IT industries, since 1965, valued at a minimum of a billion dollars each. To update this study, CISE funded CSTB to identify recent IT industries that have reached the billion dollar mark; develop a brief report that highlights the updated figures; and summarize results-to-date of IT research, including the nature and successes of U.S. research partnerships among government, industry, and universities, and the economic payoffs of these research investments. The report, *Continuing Innovation in Information Technology*,[4] was published in 2012.
- In FY 2011 through FY 2012, CISE supported several community activities to assess future research and infrastructure needs. A workshop organized by the CISE-funded Computing Community Consortium (CCC) resulted in a report, *Science, Engineering, and Education of Sustainability: The Role of Information Sciences and Engineering*.[5] This report defined a vision for fundamental research at the intersection of sustainability and IT. A subsequent study by the CSTB, *Computing Research for Sustainability*[6] specifies a framework for how innovation in computing will be essential to finding real world solutions to sustainability challenges, such as electricity production and delivery, global food production, and environmental adaptation. Separately, a CSTB study, *The Future of Computing*

[3] www.nap.edu/catalog.php?record_id=12174
[4] www.nap.edu/catalog.php?record_id=13427
[5] http://cra.org/ccc/docs/RISES_Workshop_Final_Report-5-10-2011.pdf
[6] www.nap.edu/catalog.php?record_id=13415

Performance: Game Over or Next Level?[7], together with a CCC white paper, *21st Century Computer Architecture*[8], outline the need for advances in computer architecture research. Additionally in FY 2013, the CCC collected community white papers articulating the potential needs and payoff for additional investments in mid-scale infrastructure for computing research.[9]

- CISE also funded three studies at the NAS CSTB that are currently ongoing and have the potential to influence the development of CISE programs in FY 2015.
  - *A Primer on Cybersecurity: Leveraging Two Decades of National Academies Work*: will examine what is known about effective technical and nontechnical approaches, the state of art and open challenges, why relatively little progress has been made in cybersecurity despite the recommendations of many reports from the Academies and elsewhere, and potential policy responses.
  - *Continuing Innovation in Information Technology: A Workshop:* will conduct a public workshop that highlights additional examples of the impacts of computing research using the framework established in the "tiretracks" figure published in *CSTB's 2012 report Continuing Innovation in Information Technology* and explore further uses of the figure and framework.
  - *Toward 21st-Century Cyber-Physical Systems Education*: will conduct a study on the current and future needs in education for cyber-physical systems (CPS) and articulate a vision for a 2lst century CPS-capable U.S. workforce.

The Performance chapter provides details regarding the periodic reviews of programs and portfolios of programs by external Committees of Visitors and directorate Advisory Committees. Please see this chapter for additional information.

Number of People Involved in CISE Activities

	FY 2013 Actual Estimate	FY 2014 Estimate	FY 2015 Estimate
Senior Researchers	6,652	6,900	6,900
Other Professionals	1,186	1,200	1,200
Postdoctorates	475	500	500
Graduate Students	6,609	6,900	6,900
Undergraduate Students	2,305	2,400	2,400
Total Number of People	**17,227**	**17,900**	**17,900**

[7] www.nap.edu/openbook.php?record_id=12980
[8] http://cra.org/ccc/docs/init/21stcenturyarchitecturewhitepaper.pdf
[9] http://cra.org/ccc/visioning/visioning-activities/mid-scale-infrastructure-investments-for-computing-research

DIVISION OF ADVANCED CYBERINFRASTRUCTURE (ACI) $212,287,000
$0 / 0.0%

ACI Funding

(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Change Over FY 2014 Estimate Percent
Total, ACI	**$207.59**	**$212.29**	**$212.29**	**-**	**-**
Research	**95.35**	**96.19**	**94.15**	**-2.04**	**-2.1%**
CAREER	3.96	3.47	3.47	-	-
Education	**6.70**	**5.80**	**5.84**	**0.04**	**0.7%**
Infrastructure	**105.53**	**110.30**	**112.30**	**2.00**	**1.8%**
Networking and Computational Resources	105.53	110.30	112.30	2.00	1.8%

Totals may not add due to rounding.

ACI supports the advance of science and engineering research and education by exploring, developing, creating, and supporting secure, advanced, global cyberinfrastructure (CI). ACI partners with NSF directorates and offices to support research and development across the entire range of cyberinfrastructure activities, including acquisition, integration, coordination, and operations associated with data, networking, computation, software, and the development of CDS&E tools and expertise. ACI focuses on the development of these resources and capabilities, as well as on the expertise to conduct next generation science and engineering, in order to better address complex and multidisciplinary discovery, prediction, and innovation. ACI also provides computational support to more than 8,000 faculty and researchers and supports international activities in networking, software, data, and computation, including connectivity to major international resources and scientific instruments.

In general, 41 percent of the ACI portfolio is available for new research grants and 59 percent is available for continuing grants.

Approximately 60 percent of ACI's budget is used to support individuals and small groups of researchers, while about 40 percent of the budget goes to support larger cyberinfrastructure consortia, including the eXtreme Digital (XD) services program, the Blue Waters Petascale Computing Program, and mid-scale pilots and prototypes toward an integrated data infrastructure.

FY 2015 Summary

All funding decreases/increases represent change over the FY 2014 Estimate.

Research

- ACI continues support for early-career researchers through investments in the CAREER program ($3.47 million, equal to the FY 2014 Estimate).
- In partnership with ENG, SBE, and other CISE divisions, ACI will invest in the CRISP program, focusing on the multi-disciplinary challenges related to large-scale resilient research cyberinfrastructure (+$3.0 million to a total of $3.0 million).
- ACI will continue to provide leadership for CIF21, including developing coordinated CIF21 programs and solicitations and identifying common approaches for a scalable, comprehensive cyberinfrastructure. ACI provides significant funding for multidisciplinary data and software

programs, such as DIBBs, SI^2, EarthCube and CDS&E. Research and infrastructure funding for CIF21 is reduced in FY 2015 (-$5.0 million for a total of $60.0 million).

- ACI will support Clean Energy Technology ($3.50 million, equal to the FY 2014 estimate), which includes investments in the NSF-wide SEES portfolio ($2.50 million, equal to the FY 2014 Estimate). ACI investments will focus on the exploration of new forms of cyberinfrastructure to advance sustainability science, engineering, and education.

Education

- Along with the other CISE divisions, ACI will increase its current level of investment in Research Experiences for Undergraduates (REU) supplements and sites (+$40,000 to a total of $1.34 million).
- ACI will maintain support for the CIF21 emphasis area within the NSF Research Traineeship program ($3.0 million, equal to the FY 2014 Estimate).

Infrastructure

- Advanced, secure networking is increasingly important to support research collaborations and facilitate greater data sharing across disciplines. In partnership with EHR, ENG, MPS, SBE, and other CISE divisions, ACI continues to participate in the SaTC program through the transition to practice option (-$1.0 million to a total of $3.0 million).
- The advanced computing investment (+$1.0 million to a total of $83.0 million) includes the combined services and resources of XD, operating a virtual computational environment for more than 8,000 scientists and engineers. Within XD, support for several resources will end in FY 2015, with the deployment of a new resource, Comet, more suitable for both high throughput and data-intensive computing (see Appendix A for more information on the HPC portfolio). The Blue Waters sustained petascale resource will be in its second full year of operations in FY 2015. It allows computational scientists in a variety of domains, such as molecular dynamics, protein folding, cosmology, climate modeling, and earthquake modeling, to retain international competitiveness by addressing some of the most difficult research problems known today.

DIVISION OF COMPUTING AND COMMUNICATION FOUNDATIONS (CCF) **$185,185,000**
$0 / 0.0%

CCF Funding
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Percent
Total, CCF	**$178.02**	**$185.19**	**$185.19**	**-**	**-**
Research	**168.32**	**175.30**	**175.10**	**-0.20**	**-0.1%**
CAREER	15.11	13.20	13.37	0.17	1.3%
Centers Funding (total)	8.32	12.66	10.00	-2.66	-21.0%
Team for Research in Ubiquitous Secure	3.32	2.66	-	-2.66	-100.0%
STC: Science of Information	5.00	5.00	5.00	-	-
STC: Center for Brains, Minds and Machines: the Science and the Technology of Intelligence	-	5.00	5.00	-	-
Education	**9.10**	**9.29**	**9.49**	**0.20**	**2.2%**
Infrastructure	**0.60**	**0.60**	**0.60**	**-**	**-**
National Nanotechnology Infrastructure Network	0.60	0.60	0.60	-	-

Totals may not add due to rounding.

CCF supports research and education activities that explore the foundations and limits of computation, communication, and information; advance algorithmic knowledge for research areas both within and outside computer science; and advance software, hardware, and computer system design. CCF's research investments support advances in the design and analysis of algorithms; computational complexity, theoretical and experimental studies of algorithms and their resource requirements; and formal models of computation. These research investments include models for parallel, distributed, and heterogeneous multi-core machines. CCF invests in research addressing the theoretical underpinnings and enabling technologies for information acquisition, transmission, and processing in communication and information networks, such as sensor, wireless, multimedia, and biological networks. CCF investments advance the design, verification, evaluation, and utilization of computing hardware and software through new theories and high-leverage tools that focus on performance, correctness, usability, dependability, reliability, and scalability. CCF also invests in research that explores the potential impact of emerging technologies on computation and communication, including nanotechnology, biotechnology, and quantum devices and systems.

In general, 75 percent of the CCF portfolio is available for new research grants and 25 percent is available for continuing grants.

FY 2015 Summary

All funding decreases/increases represent change over the FY 2014 Estimate.

Research

- CCF continues support for early-career researchers through increased investments in the CAREER program (+$170,000 to a total of $13.37 million).
- CCF maintains the current investment level of $4.0 million for the Exploiting Parallelism and Scalability program (XPS). CCF supports development of new foundational principles and cross-layer approaches that integrate both hardware and software through new programming languages,

models, algorithms, compilers, runtime systems, and architectures. A primary goal is to achieve scalability and energy efficiency in systems while also addressing programmability, reliability, and security.

- CCF increases support for Clean Energy Technology (+$2.02 million to a total of $10.52 million), which includes support for foundational research in energy-intelligent computing. This also includes CCF investments in the NSF-wide SEES portfolio ($4.0 million, equal to the FY 2014 Estimate). In CyberSEES, CCF continues its investments in cyber-enabled sustainability, which includes novel approaches based on optimization, randomization, simulation, and inference; large-scale data management and analytics; and smart management of engineered systems. The research addresses integrative approaches to sustainable computing and information technologies across the lifecycle of design, use and reuse with the associated consumption of energy, materials, and related resources.
- CCF maintains its research investments in CIF21 ($7.50 million, equal to the FY 2014 Estimate). CCF will emphasize its investments in Big Data foundational research, including algorithms and software tools for handling large and heterogeneous data sets; randomized streaming algorithms, which are extremely efficient for Big Data; methods for validating data; and tools for extracting knowledge to enable new discoveries.
- CCF continues its participation ($13.25 million, equal to the FY 2014 Estimate) in SaTC through investments in theories, models, algorithms, architectures, languages and tools for increased security, privacy and trust, as well as in new cryptographic approaches.
- CCF supports the NSF-wide CEMMSS program at a reduced level through CPS (-$500,000 to a total of $5.25 million) and NRI ($3.50 million). This investment emphasizes development of new methods for specification and verification of software and hardware systems useful for various sectors including advanced manufacturing.
- CCF maintains its current level of investment in the cross-cutting Smart and Connected Health (SCH) initiative ($3.0 million, equal to the FY 2014 Estimate).
- CCF continues to support two STCs, *Center for Science of Information* at Purdue University and *Center for Brains, Minds and Machines: the Science and the Technology of Intelligence* at MIT ($10.0 million, equal to the FY 2014 Estimate). The STC *Team for Research in Ubiquitous Secure Technology (TRUST)* at the University of California at Berkeley will sunset in FY 2015 (-$2.66 million to a total of zero).

Education

- Along with the other CISE divisions, CCF increases its current level of investments in REU supplements and sites (+$200,000 to a total of $3.24 million).
- CCF maintains its current level of support for the STEM-C Partnerships activity ($4.0 million).
- CCF maintains support for the CIF21 emphasis area within the NSF Research Traineeship program ($1.50 million, equal to the FY 2014 Estimate).

Infrastructure

- CCF funds the National Nanotechnology Infrastructure Network, supported primarily by ENG, at a level of $600,000.

DIVISION OF COMPUTER AND NETWORK SYSTEMS (CNS) $220,403,000
$0 / 0.0%

CNS Funding
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Change Over FY 2014 Estimate Percent
Total, CNS	**$211.03**	**$220.40**	**$220.40**	**-**	**-**
Research	**165.25**	**177.26**	**176.93**	**-0.33**	**-0.2%**
CAREER	13.99	12.90	13.06	0.16	1.2%
Education	**16.01**	**13.14**	**13.47**	**0.33**	**2.5%**
Infrastructure	**29.77**	**30.00**	**30.00**	**-**	**-**
Research Resources	29.77	30.00	30.00	-	-

Totals may not add due to rounding.

CNS supports research and education activities that advance understanding of the fundamental properties of computer systems and networks; explore new ways to address the limitations of existing computer and networked systems to make better use of these technologies; and develop novel paradigms, abstractions, and tools for designing, analyzing, and building next-generation computer and networked systems that are robust, secure, and trustworthy. CNS investments in computer systems research focus on: distributed, mobile, and embedded systems; sensing and control systems; dynamically configured, multiple-component systems; and parallel systems. CNS investments in fundamental network research create new insights into the dynamics of complex networks and explore new architectures for future-generation networks and services. CNS provides scientific leadership in cybersecurity, supporting research and education activities that will ensure society's ubiquitous and distributed computing and communications systems 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 research infrastructure resources and in the development of the computing workforce of the future.

In general, 64 percent of the CNS portfolio is available for new research grants and 36 percent is available for continuing grants.

FY 2015 Summary

All funding decreases/increases represent change over the FY 2014 Estimate.

Research

- CNS continues support for early-career researchers through increased investments in the CAREER program (+$160,000 to a total of $13.06 million).
- In partnership with EHR, ENG, MPS, SBE, and the other CISE divisions, CNS continues to lead the SaTC program (-$1.0 million to a total of $42.30 million), which aligns with the national cybersecurity strategy. SaTC will invest in game-changing research in support of CNCI, develop scientific foundations, maximize research impact, and accelerate transitions to practice, in addition to addressing education and workforce issues. SaTC will fund a diverse set of collaborative projects in areas of current critical importance, such as network and cloud security, cybereconomics, and science of security.

- In partnership with BIO, ENG, MPS, and the other CISE divisions, CNS continues to support the CEMMSS initiative and its components. Through the CPS program (-$750,000 to a total of $21.50 million), CNS will support the foundational interdisciplinary research and education underlying adaptive and pervasive smart systems, and further the understanding of fundamentals arising from grand challenge applications, such as advanced manufacturing, smart grid technologies, medical devices, and transportation networks. CNS will also maintain its investments in NRI ($5.0 million).
- In partnership with ENG, SBE, and other CISE divisions, CNS continues to support the science and engineering necessary to enable advances in large-scale resilient and interdependent infrastructures through the Critical Resilient Interdependent Infrastructure Systems and Processes (CRISP) program ($3.0 million, equal to the FY 2014 Estimate).
- CNS continues to support research in wireless communication, spectrum sharing, and mobile computing together with the development of wireless testbeds ($3.0 million, equal to the FY 2014 Estimate). CNS continues its collaboration with ENG, MPS, and other CISE divisions through the Enhancing Access to the Radio Spectrum (EARS) program.
- CNS continues to support XPS ($3.0 million, equal to the FY 2014 Estimate), investing in new approaches to the traditional computer hardware and software stack in multi-core, multi-machine, cloud, data-intensive, and highly concurrent systems.
- CNS continues to support its research investments in CIF21 ($3.0 million, equal to the FY 2014 Estimate). CNS will emphasize its investments in Big Data, including novel research in combining distributed sensing with data analytics and distributed response, as well as work on large-scale data management systems.
- CNS increases support for Clean Energy Technology (+$500,000 to a total of $4.0 million), which includes investments in the NSF-wide SEES portfolio ($2.50 million, equal to the FY 2014 Estimate), through research in widespread, heterogeneous sensing and control; and new methods for addressing power, thermal and sustainability issues in the design and operation of computing systems.
- CNS participates in other cross-cutting research programs, including SCH ($3.0 million, equal to the FY 2014 Estimate) and Cyberlearning ($2.0 million, equal to the FY 2014 Estimate).

Education

- CNS maintains investments in the CISE/EHR STEM-C Partnerships, which seek to enhance computational competencies for all students ($4.0 million, equal to the FY 2014 Estimate). This program will enlarge the pool of K-14 students and teachers who develop and practice computational competencies in a variety of contexts, as well as the pool of early postsecondary students who are engaged and have the background in computing necessary to successfully pursue degrees in computing-related and computationally-intensive fields of study.
- Along with the other CISE divisions, CNS will increase support for REU sites and supplements (+$330,000 to a total of $4.38 million).
- CNS maintains support for research traineeships through the Integrative Graduate Education and Research Traineeship (IGERT) program and the CIF21 emphasis area within the NRT ($1.39 million, equal to the FY 2014 Estimate).

Infrastructure

- Through the CISE Research Infrastructure (CRI) program ($18.0 million, equal to the FY 2014 Estimate); CNS supports acquisition, enhancement, and operation of state-of-the-art computing research infrastructure enabling high-quality computing research and education in a diverse range of institutions.
- CNS maintains support for the development of world-class, mid-scale network infrastructure, facilitating the transition from GENI to NSFCloud prototypes ($12.0 million, equal to the FY 2014 Estimate) that will begin enabling novel cloud architectures addressing emerging challenges, including high-confidence systems.

DIVISION OF INFORMATION AND INTELLIGENT SYSTEMS (IIS) **$185,185,000**
+$10,000 / 0.0%

IIS Funding
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Percent
Total, IIS	**$176.23**	**$185.18**	**$185.19**	**$0.01**	**0.0%**
Research	**167.42**	**176.28**	**176.02**	**-0.26**	**-0.1%**
CAREER	17.03	14.80	14.98	0.18	1.2%
Education	**8.81**	**8.90**	**9.17**	**0.27**	**3.0%**

Totals may not add due to rounding.

IIS supports research and education to develop and apply new IT to enhance the capabilities of people and machines to create, discover, and reason by advancing their ability to represent, collect, store, organize, visualize, and communicate data and information; develop new knowledge to support people in the design and use of IT; and advance knowledge about how computational systems can perform tasks autonomously, robustly, and flexibly.

IIS research investments support the exploration of novel theories and innovative technologies that advance our understanding of the complex and increasingly coupled relationships between people and computing, promising to enhance quality of life. Investments in information integration and informatics focus on the processes and technologies involved in creating, managing, visualizing, and fusing diverse data, information, and knowledge from disparate and uncoordinated sources within a changing landscape of computing platforms, from personal devices to globally-distributed networks. IIS also invests in research on artificial intelligence, computer vision, natural language, robotics, machine learning, computational neuroscience, cognitive science, and areas leading to the computational understanding and modeling of intelligence in complex, realistic contexts. These investments aim to revolutionize understanding of brain functions.

In general, 67 percent of the IIS portfolio is available for new research grants and 33 percent is available for continuing grants.

FY 2015 Summary

All funding decreases/increases represent change over the FY 2014 Estimate.

Research

- IIS continues support for early-career researchers through increased investments in the CAREER program (+$180,000 to a total of $14.98 million).
- IIS will increase its investments in Cognitive Science and Neuroscience (+$2.15 million to a total of $5.65 million) by supporting the integration of computational models across multiple scales, from the atomic to the behavioral. This research aims to accelerate the formulation of an integrative, quantitative, and predictive theory of brain and nervous system function.
- Partnering with three other federal agencies (NASA, NIH, and USDA) and three other NSF directorates (ENG, SBE, and EHR), IIS will continue to lead the NRI program (-$1.0 million to a total of $11.0 million). IIS will focus on fundamental research in robotics as a critical underpinning of CEMMSS, which includes advanced sensing, control, and power sources; integrated problem-

solving architectures and decision algorithms; and safe and soft structures. NRI focuses on human-centered research in developing service robots, requiring significant advances in human-robot interaction, including safety standards for robots that work beside or cooperatively with humans and recovery from errors in robot interpretation and action. Application domains include robots as co-workers in advanced manufacturing environments, aides supporting emergency responders in the field, and service robots assisting the elderly to live independently.

- As part of its CEMMSS investment, IIS also supports the CPS program (-$1.25 million to a total of $3.75 million). IIS invests in basic research in smart systems with embedded computational intelligence that can sense, adapt, and react thereby enabling new functionalities and providing next-generation products and services in various sectors, including advanced manufacturing.
- IIS continues to invest in CIF21 ($9.50 million, equal to the FY 2014 Estimate). In partnership with ENG, EHR, and SBE and other CISE divisions, IIS will lead Big Data research activities, addressing challenges in data management, data analytics, and scientific discovery processes. IIS research will focus on novel theoretical analysis or experimental evaluation of these techniques and methodologies. Topics include data mining methods, and data- and information-fusion techniques, machine learning, predictive modeling, and automated discovery of phenomena and causality in data. IIS will provide integration and transition supplement awards for joint work between big data core technology research and exploratory domain application projects and/or joint industry collaboration projects.
- IIS continues to lead the Cyberlearning program ($10.0 million, equal to the FY 2014 Estimate) jointly with EHR and SBE. This research will integrate advances in technology with advances in the ways people learn; resolve how to more effectively use technology for promoting learning; and design new technologies for integration in learning environments and evaluate their use. Interdisciplinary teams will study the deluge of data produced from new on-line teaching paradigms, such as that from massive open on-line courses, to better understand learning mechanisms and enable productive, personalized, and customized education.
- In partnership with NIH, NSF directorates ENG and SBE, and CISE divisions CCF and CNS, IIS will continue to lead the SCH program ($9.0 million, equal to the FY 2014 Estimate). IIS will pursue improvements in safe, effective, efficient, and patient-centered proactive and predictive health and wellness technologies through innovations in computer and information science and engineering. The program addresses changing age demographics with investments in assistive cyber-physical engineered systems that are embedded in the local environment and are easily accessed and used.
- IIS continues to support the NSF-wide SaTC program (-$1.0 million to a total of $8.45 million) through research in the data science foundations of a secure and trustworthy cyberspace and privacy-protecting mechanisms for heterogeneous fused data.
- IIS continues to invest in Clean Energy Technology ($2.50 million, equal to the FY 2014 estimate), which includes investments in the NSF-wide SEES portfolio ($2.0 million, equal to the FY 2014 Estimate). CyberSEES investments focus on the information processing dimensions of energy utilization and pursue breakthroughs needed in optimization, modeling, simulation, and inference.

Education

- With all CISE divisions, IIS increases support for REU sites and supplements (+$270,000 to a total of $3.92 million).
- IIS maintains investments in the CISE/EHR STEM-C Partnerships activity ($4.0 million, equal to the FY 2014 Estimate), which seeks to enhance computational competencies for students and teachers.
- IIS maintains support for the CIF21 emphasis area within the NSF Research Traineeship program ($500,000).

DIVISION OF INFORMATION TECHNOLOGY RESEARCH (ITR) $90,290,000
-$660,000 / -0.7%

ITR Funding
(Dollars in Millions)

	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request	Change Over FY 2014 Estimate Amount	Change Over FY 2014 Estimate Percent
Total, ITR	**$85.25**	**$90.95**	**$90.29**	**-$0.66**	**-0.7%**
Research	**76.51**	**82.20**	**78.84**	**-3.36**	**-4.1%**
CAREER	0.03	-	-	-	N/A
SLC: Pittsburgh Science of Learning Center - LearnLab	2.25	2.03	-	-2.03	-100.0%
Education	**0.12**	**0.75**	**1.45**	**0.70**	**93.3%**
Infrastructure	**8.63**	**8.00**	**10.00**	**2.00**	**25.0%**
Research Resources	8.63	8.00	10.00	2.00	25.0%

Totals may not add due to rounding.

ITR provides support for transformative explorations in computer and information science and engineering research, infrastructure, and related education activities, emphasizing the funding of high-risk, multi-investigator, and multidisciplinary projects.

In general, 62 percent of the ITR portfolio is available for new research grants and 38 percent is available for continuing grants.

FY 2015 Summary

All funding decreases/increases represent change over the FY 2014 Estimate.

- Through I-Corps (+$2.50 million to a total of $10.0 million), ITR will invest in NSF-funded researchers who will receive additional support – in the form of mentoring and funding – to accelerate innovation and transfer of knowledge from lab to practice. ITR will also invest in two additional I-Corps subcomponents that were initiated in FY 2013 – Sites and Nodes – to further build, utilize, and sustain a national innovation ecosystem that continues to augment the development of technologies, products, and processes that benefit the Nation.
- Leveraging previous investments in GENI, ITR will invest in US Ignite ($3.0 million), an effort to promote U.S. leadership in developing gigabit public sector applications and services for ultra-fast broadband and software-defined networks. As part of US Ignite, ITR will invest in foundational wireline, wireless, cloud computing, security, and distributed systems research and experimentation, as well as gigabit application development.
- In collaboration with ENG, CISE will continue to support innovative partnerships and collaborations between universities and industries, in part through the Industry/University Cooperative Research Centers (I/UCRC) program, which will continue to establish centers that partner industry with university research efforts ($8.0 million).
- ITR will maintain its investments in the Expeditions in Computing program ($12.0 million, equal to the FY 2014 Estimate). This program identifies projects with transformative research agendas that promise to accelerate discovery at the frontiers of computing and communication. It will also continue to encourage researchers to come together within or across departments and/or institutions to

identify the compelling ideas that promise transformations in computing and information sciences for many years to come.

- ITR will continue to support development and deployment of wireless testbeds (-$250,000 to a total of $4.0 million).
- ITR will invest in multi-disciplinary research networks, including support for the Science Across Virtual Institutes (SAVI) activity ($2.0 million, equal to the FY 2014 Estimate). These research networks will provide opportunities to develop collaborations in areas of emerging interest to computer and information science and engineering, including international partnerships.
- ITR will continue to provide support for emerging and urgent high-priority areas of potentially transformative research through various award mechanisms, such as Early-concept Grants for Exploratory Research (EAGERs) and Grants for Rapid Response Research (RAPIDs), and through co-funding of awards with other NSF directorates to pursue important emerging areas.
- LearnLab, formerly known as the Pittsburgh Science of Learning Center (SLC) for Robust Learning, will sunset in FY 2015 (-$2.03 million to a total of zero).

Education

- ITR will increase support for the CIF21 emphasis area within the NSF Research Traineeship program (+$700,000 to a total of $1.20 million). This program is designed to encourage the development of new, potentially transformative, and scalable models of STEM graduate training that ensure graduate students develop the skills, knowledge, and competencies needed to pursue a range of careers within and outside academe.

Infrastructure

- As part of US Ignite, ITR will maintain its mid-scale infrastructure investments in GENI, at a funding level of $4.0 million, research-enabling and/or integrating additional U.S. campuses, regional and research backbone networks, commercial equipment, and cities across the nation to create a unique at-scale infrastructure available for future network research and infrastructure experimentation.
- Transitioning investments in GENI, ITR will extend virtualization beyond the network to large-scale, interconnected computing resources by developing mid-scale prototypes for an NSFCloud research infrastructure, enabling future cloud and distributed computing experimentation not otherwise possible ($6.0 million).

APPENDIX A – HIGH PERFORMANCE COMPUTING PORTFOLIO

High Performance Computing Funding
(Dollars in Millions)

	Prior Years	FY 2013 Actual	FY 2014 Estimate	FY 2015 Request
Petascale (Track 1)	$269.52	$30.12	$30.40	$30.00
Innovative HPC Program	201.06	12.32	32.30	30.00
Teragrid - Phase III (XD)	255.71	27.44	19.30	23.00
Total	**$726.29**	**$69.88**	**$82.00**	**$83.00**

Totals may not add due to rounding.

NSF has been a leader in the use of High Performance Computing (HPC) to advance discovery for almost four decades. As a result of continuous rapid changes in computing and related technologies, coupled with the exponential growth and complexity of data for the science, engineering, and education enterprise, NSF has created a new vision and strategy towards Advanced Computing Infrastructure (ACI), which will expand NSF's leadership role in science and engineering. This coordinated NSF-wide strategy, which is a key component of the Cyberinfrastructure Framework for 21st Century Science, Engineering, and Education (CIF21) framework, seeks to position and support the entire spectrum of NSF-funded communities at the cutting edge of advanced computing technologies, hardware, and software. It also aims to promote a more complementary, comprehensive, and balanced portfolio of advanced computing infrastructure and programs for research and education. The strategy enables multidisciplinary computational- and data-enabled science and engineering that supports all science, engineering, and education communities. This shift is consistent with the recommendations of a 2010 review and 2012 follow-up review[10] of the Federal Networking and Information Technology R&D (NITRD) program by the President's Council of Advisors on Science and Technology (PCAST).

PETASCALE COMPUTING (TRACK 1) – BLUE WATERS

Description

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC) is now providing the capability for researchers to tackle much larger and more complex research challenges than previously possible. This was accomplished by acquiring, deploying, and operating a sustained petascale leadership-class, high-performance, computational resource known as Blue Waters. It is important to note that this investment complements the Department of Energy (DOE) Office of Science's program on computing hardware, which focuses on peak petascale performance. In contrast, Blue Waters provides sustained petascale performance. Blue Waters also complements the broad set of resources provided through the eXtreme Digital (XD) environment, which provide advanced computational support to the U.S. science and engineering community. Blue Waters focuses on a small set of the largest and most computational-intense scientific advances while XD delivers a broader set of capabilities to a much larger community. (For more information on XD, see the discussion on TeraGrid Phase III that is included within this appendix.) A Blue Waters Early Science System was deployed at the University of Illinois at Urbana-Champaign in FY 2012. The full system was operational in December 2012, and the archival storage availability came online in March 2013. It is operated by NCSA and its partners in the Great Lakes Consortium for Petascale Computing (GLC).

[10] *Designing a Digital Future: Federally Funded Research and Development Networking and Information Technology,* President's Council of Advisors on Science and Technology, January 2013.

The Blue Waters project includes education and outreach programs that target pre-college, undergraduate, graduate, and post-graduate levels. A Virtual School of Computational Science and Engineering was established to create courses that focus on petascale computing and petascale-enabled science and engineering. The Virtual School is exploring new instructional technologies and creating courses, curricula, and certificate programs tailored to science and engineering students. It has also sponsored workshops, conferences, summer schools, and seminars.

The project includes an annual series of workshops targeted at the developers of simulation packages and aspiring application developers. In addition, the project includes two industrial partnership activities. The Industry Partners in Petascale Engagement (IPIPE) program provides industrial partners with a first look at the technological and scientific developments that flow from the petascale program. The Independent Software Vendor Application Scalability Forum promotes collaborations among consortium members, independent software vendors, and the industrial end-user community.

The broader impacts of this award include: provisioning of unique infrastructure for research and education; extensive efforts accelerating education and training in the use of high-performance computation in science; training in petascale computing techniques; promoting an exchange of information between academia and industry about the applications of petascale computing; and broadening participation in computational science through NCSA's Girls Engaged in Mathematics and Science (GEMS) program. GEMS is designed to encourage middle-school girls to consider mathematics- and science-oriented careers.

Current Status

In late September 2011, Cray was selected as the vendor for the Blue Waters project. A Blue Waters Early Science System, representing about 15 percent of the overall capacity of the system, became operational in March 2012 and was used by 12 different research teams. Following system testing and acceptance in December 2012, and acceptance of the NCSA archival system in March 2013, the Blue Waters project entered a five-year operations phase. Support for the first six months of operations was provided in the acquisition and deployment award while a renewal proposal from UIUC for operations was submitted in FY 2012 and approved in FY 2013. The renewal award covers the remaining operational phase, through mid-FY 2018.

The Blue Waters education and outreach projects are ongoing, with components on undergraduate education, graduate education, training workshops, and outreach. A December 2011 workshop provided scientists and engineers with knowledge and expertise to develop applications for Blue Waters and other petascale computers. Annual extreme scale workshops are held jointly with the Extreme Science and Engineering Discovery Environment (XSEDE) project. The Blue Waters team hosts summer workshops and has created and offered courses through a virtual school of computational science and engineering. Partnering with the Shodor Foundation, a nonprofit national resource for computational science education, the Blue Waters project has offered undergraduate course materials and internships.

Science and engineering research and education activities enabled by Blue Waters

Now that Blue Waters is fully operational, it is enabling investigators across the country to conduct innovative research demanding petascale capabilities. In particular, previously awarded allocations of time on Blue Waters have begun for approximately 30 teams. These research teams were awarded time on Blue Waters through the Petascale Computing Resource Allocations (PRAC) process. More than 120 requests for usage were submitted across a wide spectrum of research areas. The research topics include: complex biological behavior in fluctuating environments; the electronic properties of strongly correlated systems; the properties of hydrogen and hydrogen-helium mixtures in astrophysically relevant conditions; the electronic and magnetic structures of transition metal compounds; the molecular dynamics responsible for the properties of liquid water; and the propagation of seismic energy through a detailed structural

model of Southern California together with the predicting of ground motion and the modeling of the response of buildings and other structures. Other allocations address testing hypotheses about the role of cloud processes and ocean mesoscale eddy mixing in the dynamics of climate and improving climate models; the formation of the first galaxies; turbulent stellar hydrodynamics; binary black hole and neutron star systems as sources of gamma ray bursts; and other intense radiation phenomena, contagion, and particle physics.

Early research results are promising. For example in May 2013, the cover of the journal *Nature* displayed the complete 64-million HIV capsid structure revealed by simulations conducted on Blue Waters.[11] The accompanying report detailed the findings, which combined experimental data with unprecedented simulations on Blue Waters, to reveal the structure for the first time. The capsid has become an attractive target for the development of new antiretroviral drugs. A new PRAC allocation solicitation was issued in December 2013 with approximately 15-20 awards anticipated by the end of FY 2014. Now that the project has entered full operations, it has issued calls for educational allocations directly involving students, including the Blue Waters undergraduate student internship program and the Blue Waters Fellowship program. Awards will begin in FY 2014.

Management and Oversight

NSF Structure: The project is overseen by CISE/ACI program staff and a grants officer from the Division of Grants and Agreements (DGA). These NSF staff members receive strategic advice from the CIF21 Steering Committee, which includes assistant directors (ADs) and office directors (ODs) from the various research directorates and offices. Advice from the Office of General Counsel (OGC) is sought, as necessary.

External Structure: During the development and acquisition phase of this project UIUC oversaw work by a number of sub-awardees, conducted software development, and assisted competitively selected research groups to prepare to use the Blue Waters system. The primary sub-awardee, Cray, is responsible for maintenance of the hardware, system software, and main program development tools. Other sub-awardees worked on extreme-scale parallel algorithm and method development, the engagement of applications groups, scalable performance tools, undergraduate training, and broadening the participation of underrepresented groups in high-performance computing. During the operational phase, the project team is advised by a Petascale Executive User Committee composed of representatives from research teams with Blue Waters allocations, industry scientists pursuing petascale applications, and other extreme-scale application experts.

Risks: Now that Blue Waters is operational, the major risks are retired. The NSB will receive updates on any major change in risk assessment.

Reviews: The project was selected through a competitive merit review process in 2007. An external panel of experts, selected by NSF, periodically reviews the progress of the project including project management, risk management, hardware and software development, and the provision of advanced user support to research groups receiving provisional resource allocations on the Blue Waters system. One of the important roles of this external review panel is to analyze the awardee's assessments of the deliverables from its sub-awardees, together with the awardee's and sub-awardees' plans for remedial action, when necessary, and to provide NSF with advice on whether these assessments and plans are reasonable. These external reviews were conducted in February 2008, April 2008, October 2008, April 2009, July 2009, December 2009, April 2010, September 2010, December 2010, February 2011, May 2011, September 2011, March 2012, August 2012, December 2012, and July 2013. The next review is planned for May 2014.

[11] www.nature.com/nature/journal/v497/n7451/index.html

INNOVATIVE HPC PROGRAM

Description

Using lessons learned during the execution of the HPC Track 2 program and informed by the NSF Advisory Committee for Cyberinfrastructure's (ACCI) High Performance Computing task force, the HPC Track 2 program was renamed Innovative HPC in 2011. Innovative HPC awards are made in the context of the XD services program (described below). While the Petascale Computing (Track 1) system is targeted to provide sustained petascale performance, the Innovative HPC systems provide, at most, petascale peak performance. Each system is capable of supporting hundreds to thousands of researchers (over the course of a year) conducting leading-edge science and engineering. The portfolio of systems supported by the Innovative HPC program is intended to be technically diverse, reflecting changing and growing use of computation in both the research and education process. NSF's support complements and extends campus and regional research cyberinfrastructures.

A direct relationship exists between the Innovative HPC awards and the XD program, which is described below. Several systems are currently serving as allocable resources within XD. Initially, Innovative HPC awards were generally made as two parts: a) an acquisition component and associated funding, and b) an operations and maintenance component and associated funding. More recent awards in the Innovative HPC program (including FutureGrid, Gordon, and Keeneland) did not separate these components because of the experimental nature of the systems. When an award was made, funding was provided to the institution, which issued sub-awards to vendors for acquisitions as necessary. Once the system has passed the acceptance process, vendors receive final payment for the system. After the system has been fully tested, it becomes an XD resource and the institution becomes an XD resource provider and has access to the operations and maintenance funding component of the award.

Beginning with the FY 2011 solicitation, *High Performance System Acquisition: Enhancing the Petascale Computing Environment for Science and Engineering*, which was based on feedback from the scientific and engineering community, a more sustained approach to core HPC services was initiated. This approach provides a longer time horizon for funding HPC providers in recognition of the value and time required for building and retaining staff skilled in interdisciplinary computational science. Thus, an eight to ten year award horizon is envisioned for a core HPC provider. This timeline begins with an acquisition award, which allows for the possibility of a renewal acquisition award four years after the original award. In addition to the acquisition awards, accompanying operations and maintenance (O&M) awards are planned.

Current Status

Machines and facilities that have been operational in the Innovative HPC program include Stampede, Blacklight, FutureGrid, Gordon, Keeneland, Kraken, Lonestar, Longhorn, and Trestles. Blacklight and Longhorn are no longer part of the current NSF HPC portfolio, although Blacklight is still supported by the Data cluster in CISE/ACI. NSF support for FutureGrid, Gordon, Keeneland, Kraken, Lonestar, and Trestles is scheduled to end by 2015.

A new solicitation for innovative computational and data resources was developed and issued in the spring of 2013. The CISE/ACI HPC and Data clusters jointly sponsored the solicitation, with funding from both sources. The proposals submitted were reviewed in the summer of 2013 and two awards were made late in FY 2013. The first award, Wrangler, will be online in January 2015 at the University of Texas at Austin, and is supported by the CISE/ACI Data cluster. Upon its deployment, Wrangler will be the most powerful data analysis system allocated in XD, with 10 petabytes (PB) of replicated, secure, high performance data storage. It will consist of 3,000 embedded processing cores for data analysis; 120 Intel Haswell-based servers for data access and embedded analytics; and a large-scale flash storage tier for analytics, with bandwidth of 1 terabyte per second (TB/s) and 275 million Input/Output Operations

Per Second (IOPS). The system will provide flexible support for a wide range of software stacks, including Hadoop and relational data, as well as integrate with Globus Online services for rapid and reliable data transfer and sharing.

The second award, Comet, also scheduled to come online in January 2015 at the University of California at San Diego, is supported by the HPC program. It is designed to be part of an emerging cyberinfrastructure for the "long tail of science," which encompasses the idea that a large number of modestly sized computationally-based research projects still represents, in aggregate, a tremendous amount of research and scientific impact. Notably, as a resource that is responsive to the "long tail of science," Comet is particularly well-suited for science gateway use. Its heterogeneous configuration will support not only complex simulations, but also advanced analytics and visualization of output.

A new solicitation was issued in FY 2014 (*High Performance Computing System Acquisition: Continuing the Building of a More Inclusive Computing Environment for Science and Engineering*) to continue the intent of the FY 2013 solicitation, which aims to add complementary computational resources to the NSF portfolio. Compared to earlier solicitations, the FY 2013 and FY 2014 solicitations were designed to broaden the spectrum of the program by exploring new and creative approaches to delivering innovative computational resources to an increasingly diverse community and portfolio of scientific research and education projects. The goal is to include new communities with needs that are different than the more traditional HPC users, but which would benefit from advanced computational capabilities at the national level. It is expected that a significant upgrade of NSF resources will occur in FY 2015, with the deployment of both Comet and Wrangler, as well as the planned upgrade of Stampede. These additions replace the resources mentioned above that will no longer be supported, primarily Kraken and Gordon.

Science and engineering research and education activities enabled by Innovative HPC

- The complete spectrum of scientific research is supported, including: climate and weather modeling, economics, cosmology and astrophysics, geosciences, physics, chemistry, biology and medicine, earthquake engineering, and mechanical engineering.
- Innovative HPC will enable world-leading transformative advances in science and engineering research, in the integration of research and education, and in broadening participation in science and engineering by underrepresented groups. These advances will be accomplished by providing researchers and educators with usable access to computational resources beyond those typically available on most campuses, together with the interfaces, consulting support, and training necessary to facilitate their use.
- Through the unifying XD framework and services, Innovative HPC will enable researchers to manipulate extremely large amounts of digital information from simulation, sensors, and experiments, and add needed capabilities in remote visualization, an increasingly important analysis tool for modern science and engineering.
- Outreach and training critical to reducing the barriers to the use of HPC systems by the research and education community will be provided by engaging research universities and foundations. Innovative HPC will incorporate new computational technologies and new approaches to software and data management, together with the expertise to enable researchers and students to complement theory and experiment with an equal emphasis in computation.

Management and Oversight

NSF Structure: CISE/ACI's program directors (PDs) provide direct oversight during both the acquisition and operations phase. Formal reporting consists of quarterly and annual reports, which are reviewed by the PDs. The PDs also hold bi-weekly teleconferences with the awardees.

External Structure: Each Innovative HPC award is managed under a cooperative agreement. Each

awardee is responsible for the satisfactory completion of milestones in order for the spending authorization to be raised. Progress is evaluated by annual reviews and the NSF PDs.

Each project has a detailed management plan in place. Each cooperative agreement includes the management structure, milestones, spending authorization levels, and review schedule.

Risks: Any activity of this nature, and at this scale, comes with a certain element of risk. The review process, conducted prior to award, reviews and analyzes the risks as presented in the proposal and identifies any additional risk that should be considered. The awards are experimental in nature, therefore, encompassing high-risk, high-reward scenarios. The award process requires that risks be identified and analyzed, and that a mitigation plan be created and followed. One of the activities of the periodic NSF external reviews, conducted by a panel of experts, is to revisit and assess the risk situation and make recommendations as deemed necessary. Risks that are no longer applicable are retired. New risks may be added, or degree of risk promoted or demoted as necessary, all of which is documented in a risk register.

Reviews: Annual reviews are performed as part of the XD review. Semi-annual reviews are performed as part of the acquisition phase. The reviews are managed by NSF PDs. The reviewers' backgrounds include scientific research, project management, large-scale systems acquisitions and operations, and a familiarity with projects funded by NSF, as well as other federal agencies. To the extent possible, continuity through this series of reviews is provided by using the same set of reviewers.

The most recent awards in the Innovative HPC program are described below.

Comet – Gateways to Discovery: Cyberinfrastructure for the Long Tail of Science, at the University of California San Diego

Description

- The Comet project at the University of California San Diego is designed to efficiently deliver significant computing capacity (two petaflops) for the 98 percent of research that requires fewer than 1,000 simultaneous and tightly coupled cores to be conducted. This project is the first NSF-supported virtualized HPC cluster at the national level.
- Comet will provide novel approaches for resource allocation, scheduling, and user support; queues with quicker response for high-throughput computing, medium-term storage allocations, and easy access to virtualized environments with customizable software stacks. Comet will also provide dedicated allocations of physical/virtual machines, improved support for science gateways, and bandwidth reservations on high-speed networks.
- Comet's heterogeneous architecture includes specialized nodes with a variety of capabilities, such as visualization, large memory, and flash memory for local node disk.
- Comet is scheduled to be deployed in FY 2015.

Stampede – Enabling, Enhancing and Extending Petascale Computing for Science and Engineering, at the University of Texas at Austin

Description

- The Stampede project at the University of Texas at Austin delivered a new system for allocation of NSF XD cyberinfrastructure services in January 2013. Stampede replaces a previous system that was developed from an award in FY 2007.
- The new resource and accompanying services target science and engineering researchers using both advanced computational methods and emerging data-intensive approaches.

- The system has boosted XD resources to nearly twice their previous capacity. It is also providing researchers with early access to a potentially transformative new approach to performance via Intel Many Integrated Core (MIC) processors, which were accepted in August 2013. The implementation of second generation MIC processors is planned in late FY 2015.
- Per the cooperative agreement and award terms, and depending on project reviews, a renewal submission to replace/upgrade Stampede may be permitted.

TERAGRID PHASE III: EXTREME DIGITAL (XD)

Description

- XD, successor to the TeraGrid program, creates and maintains an advanced, nationally distributed, open cyberinfrastructure comprised of shared user and management services, supercomputing, storage, analysis, visualization systems, data services, and science gateways connected by high-bandwidth networks, integrated by coordinated policies and operations, and supported by computing and technology experts.
- XD enables and supports leading-edge scientific discovery and promotes science and technology education.
- XD has taken a significant step forward by encouraging innovation in the design and implementation of an effective, efficient, increasingly virtualized approach to the provision of high-end digital services – extreme digital services – while ensuring that the infrastructure continues to deliver high-quality access for the many researchers and educators that use it in their work.

Current Status

Three awards are currently active within the XD program. Two smaller awards, Technology Audit (TAS) and Technology Insertion (TIS), were made in FY 2010 to the University of Buffalo and UIUC, respectively. The largest award, a 5-year award of $121.13 million – known as the XSEDE project – was made to UIUC in July 2011 for the core user, educational and management services of XD. The four additional major partners in XSEDE are the University of Texas/Austin (Texas Advanced Computer Center), the University of Pittsburgh (Pittsburgh Supercomputer Center), the University of Tennessee/Knoxville (National Center for Computational Science) and the University of California/San Diego (San Diego Supercomputer Center). XSEDE had its first two annual reviews in June 2012 and 2013 at NSF. The project will be reviewed comprehensively in September 2014; and based on this review, a decision will be made between renewing or re-competing the award for a second 5-year period in FY 2016.

Science and engineering research and education activities enabled by XD

- XD services enable transformative advances in science and engineering research, in the integration of research and education, and in broadening participation in science and engineering to under-represented groups. This is accomplished by providing researchers and educators with coherent and highly usable access to extreme-scale digital resources beyond those typically available on most campuses, together with the interfaces, consulting, advanced user support, and training necessary to facilitate their use.
- XD provides high-performance computing services; enables researchers to manipulate extremely large amounts of digital information from simulations, sensors, and experiments; and adds needed capabilities in remote visualization, an increasingly important analysis tool for modern science and engineering.
- XD's XSEDE project is developing tools and services that not only link users to national facilities, but also enable scientific collaborations of geographically distributed teams. In doing so, it facilitates access to digital resources and experimental testbeds within and across university campuses, as well

as government laboratories.

- The XSEDE project includes outreach and training critical to reducing the barriers to the use of advanced digital systems by the research and education communities. The project incorporates new ideas and technologies to enable researchers and students to move transparently between local and national resources, substantially lowering the barriers to effective use of cyberinfrastructure and promoting enhanced productivity.

Management and Oversight

NSF Structure:

- XD shared services consist of several interrelated parts – High-Performance Remote Visualization Service (HPRVS); Technology Audit Service (TAS); Technology Insertion Service (TIS); Coordination and Management Service (CMS); Extended Collaborative Support Service (ECSS – formerly Advanced User Support Service); and Training, Education and Outreach Service (TEOS). The last three elements constitute the XSEDE project; the other elements are implemented via separate awards. The HPRVS service is no longer active as the capability has been largely replaced within the resources under the Innovative HPC program.
- These elements are designed and implemented in a way that is consistent with sound system engineering principles, clearly tied to the user requirements of the science and engineering research community using a flexible methodology that permits the architecture to evolve in response to changing user needs and presents the individual user with a common user environment regardless of where the resources or user is located.
- The XD program is managed by CISE/ACI, informed by the ACCI and its task forces, with ongoing strategic guidance from the NSF cross-directorate CIF21 Steering Committee and Cyberinfrastructure Coordination and Leadership Group. CISE/ACI's program officers oversee the TAS, TIS, and XSEDE projects. To ensure that all stakeholders can provide project input, XSEDE, which encompasses multiple services and is the largest of the projects, has an external advisory board, a user board, and a service provider forum. CISE/ACI oversight of the XSEDE project includes participation in weekly teleconferences with senior XSEDE personnel and attending the quarterly project-wide staff meetings. Formal reporting consists of quarterly and annual reports, which are reviewed by the program officer.

External Structure: The final configuration of XD consists of an access and accompanying services component, and compute, visualization, and storage resources at a number of sites. The sites contain a range of high-performance computing platforms; large disk storage devices; computational platforms specifically tailored for remote visualization; high-bandwidth networks; a broad set of user services; and an education, outreach, and training component designed to fulfill the needs of current users of high-performance computing, as well as to broaden participation to new communities and under-represented groups in science and engineering. The composition of these sites will change in time as new resources become part of the XD family and other resources are retired. University partners may be part of XSEDE services either by providing and receiving services to the project, or by simply using the digital products being developed by XSEDE in their own local environment.