OFFICE OF CYBERINFRASTRUCTURE (OCI) $218,270,000
+$6,630,000 / 3.1%

OCI Funding

(Dollars in Millions)

	FY 2011 Actual[1]	FY 2012 Estimate	FY 2013 Request	Change Over FY 2012 Estimate Amount	Percent
Research	**$68.01**	**$87.48**	**$94.87**	**$7.39**	**8.4%**
CAREER	$3.89	$3.97	$4.16	$0.19	4.8%
Software	8.60	25.43	25.90	0.47	1.8%
Data	24.02	24.85	28.33	3.48	14.0%
Virtual Organizations	4.78	5.00	5.50	0.50	10.0%
Other Disciplinary and Interdisciplinary Research	26.72	28.23	30.98	2.75	9.7%
Education	**$11.39**	**$7.60**	**$6.40**	**-$1.20**	**-15.8%**
Infrastructure	**221.35**	**116.56**	**117.00**	**0.44**	**0.4%**
High Performance Computing (HPC)	187.91	88.31	82.00	-6.31	-7.1%
Other Networking and Computational Programs	33.44	28.25	35.00	6.75	23.9%
Total, OCI	**$300.75**	**$211.64**	**$218.27**	**$6.63**	**3.1%**

Totals may not add due to rounding.

[1] FY 2011 Actual includes $90.50 million in funds that were obligated in FY 2010, de-obligated in FY 2011 and then obligated to other projects in the OCI portfolio.

About OCI

OCI enables science and engineering research and education by creating secure, advanced, global cyberinfrastructure (CI). OCI supports research and development activities on cyberinfrastructure components as well as acquisition and operation of advanced CI. CI is used to transform data into knowledge and to develop a better understanding of complexity through simulation, prediction, and multidisciplinary collaboration.

Due to the continuing rapid change in High Performance 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). This new NSF-wide strategy will focus on a complementary, comprehensive, and balanced portfolio that will address foundational research issues in algorithms, software, parallel systems, data analytics, visualization, simulation, computational and data enabled science, future workforce issues, and hardware acquisition.

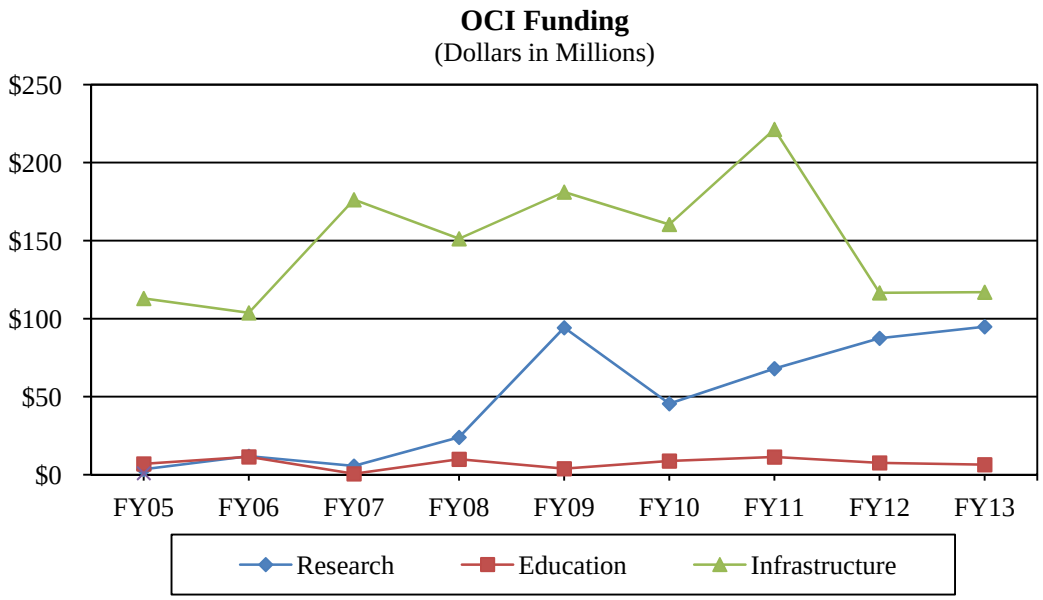

FY 2009 funding reflects both the FY 2009 omnibus appropriation and funding provided through the American Recovery and Reinvestment Act of 2009 (P.L. 111-5).

FY 2013 Summary

All funding decreases/increases represent changes over the FY 2012 Estimate.

Research

Scientific communities are being inundated by a deluge of data. Miniaturization and commoditization of computer chips and sensors, as well as the lower cost of computing and storage, contribute to large amounts of data being generated by scientific instruments and facilities. The OCI FY 2013 Request reflects the continued and growing need to support sharing of large data sets and connections among scientific communities.

- Software funding is supported at $25.90 million, an increase of $470,000 over the FY 2012 Estimate. The software program supports sustainability and extensibility of software to ensure robustness so that tools may be used across scientific communities. In FY 2013, in conjunction with the Cyberinfrastructure Framework for 21st Century Science and Engineering (CIF21) NSF-wide investment, the software program will fund new software institutes designed to help researchers extend and maintain their software.
- Data support increases by $3.48 million from the FY 2012 Estimate to a total of $28.33 million for data-enabled science, including long-term data support and infrastructure, a data life-cycle program focus (access, curation, mining, security, and management), data tools, data interoperability, data repositories, and a multi-disciplinary focus on data services, data science, and data-intensive science in conjunction with the CIF21 portfolio.
- The Virtual Organizations program increases by $500,000 to $5.50 million. Efforts in FY 2012 focus on understanding and developing new research community networks (RCNs). This includes the support of workshops and a study to understand how scientists collaborate using technology and to identify gaps in what is known. In FY 2013 efforts will be expanded in building and supporting multidisciplinary community research networks.
- Other disciplinary and interdisciplinary research funding increases by $2.75 million to a total of $30.98 million. This includes an increase of $1.50 million, to a total of $3.0 million, for Science,

Engineering, and Education for Sustainability (SEES), an increase of $750,000 to a total of $1.0 million for NSF Innovation Corps (I-Corps), and an additional $500,000 to a total of $1.0 million for INSPIRE.

Education

Support for education decreases by $1.20 million to a total of $6.40 million.

- Funding for the Cyberinfrastructure Training, Education, Advancement and Mentoring (CI-TEAM) program is eliminated (-$4.0 million) as more support is provided to NSF Fellowships for Transformative Computational Science using CyberInfrastructure (CI-TraCS) and to education through co-funding of awards with other directorates in conjunction with CIF21 activities.
- Funding for the CI-TRaCS program is increased by $2.90 million to $4.90 million to support 4 new post-doctoral researchers in their training in cyberinfrastructure. The CI-TraCS program is undergoing an evolution in FY 2013. When this evolution is complete in FY 2014, the program will more broadly support education at different levels (graduate, post-doctoral, mid-career) across the spectrum of OCI program areas and in scientific domains supported by CIF21 activities.
- OCI support for the CISE-led Computing Education for the 21st Century (CE21) program ends, for a reduction of $500,000. OCI support for this program is being eliminated to allow OCI to increase support for Research Experiences for Undergraduates (REU).
- REU site funding increases by $300,000 to $1.30 million and REU supplement funding increases by $100,000 to $200,000. OCI is committed to the support of undergraduates and to connecting smaller undergraduate institutions with larger research universities.

Infrastructure

Support for High Performance Computing decreases by $6.31 million to a total of $82.0 million.

- The Blue Waters (Track One) program is expected to be operational by the end of FY 2012. This program transitions from acquisition to operations and maintenance in FY 2013 at a funding level of $30.0 million. This is a decrease of $80.50 million from FY 2012 (including FY 2011 funds that will be obligated in FY 2012), when funds were provided for the final phase of acquisition.
- Innovative HPC funding decreases by $23.90 million to a total of $25.0 million. FY 2013 funding will support the acquisition of a single system or multiple smaller systems via a competitive process. In FY 2013, funding for operations and maintenance (O&M) is not required, as FY 2012 funds will support multiple years of O&M on several systems.
- Funding for eXtreme Digital will be $27.0 million to support ongoing commitments to these resources and services.

Support for other networking and computational programs increases by $6.75 million to $35.0 million.

- Total OCI funding for cybersecurity in FY 2013 is unchanged from FY 2012 at $21.0 million. Decreases of $1.50 million for the Comprehensive National Cybersecurity Initiative (CNCI) and of $500,000 in other cybersecurity efforts in core HPC programs are offset by an increase of $2.0 million for the OneNSF Secure and Trustworthy Cyberspace (SaTC) investment. OCI's support of SaTC, which will focus on the transition to practice perspective, is part of OCI's support for CNCI. Additional funding under CNCI supports awards for testing of new and innovative cybersecurity prototypes.
- An increase of $2.50 million will fund networking and connectivity support for components of the EarthCube project with GEO and CIF21 infrastructure projects with other directorates.
- An increase of $3.25 million to a total of $13.0 million will support the International Research Network Connections (IRNC) program to enable high-speed connections to compute facilities and NSF-supported Major Research Equipment and Facilities Construction (MREFC) projects.
- An increase of $1.0 million to a total of $6.50 million will support broadband efforts, some in conjunction with CIF21 programs.

Major Investments

OCI Major Investments

(Dollars in Millions)

Area of Investment	FY 2011 Actual	FY 2012 Estimate	FY 2013 Request	Change Over FY 2012 Estimate Amount	Change Over FY 2012 Estimate Percent
CAREER	$3.89	$3.97	$4.16	$0.19	4.8%
CEMMSS	-	0.50	1.00	0.50	100.0%
CIF21	-	23.00	32.03	9.03	39.3%
CNCI	-	8.00	6.50	-1.50	-18.8%
E^2	-	-	1.50	1.50	N/A
I-Corps	-	0.25	1.00	0.75	300.0%
INSPIRE	-	0.50	1.00	0.50	100.0%
SEES	5.61	1.50	3.00	1.50	100.0%
SaTC	-	4.00	6.00	2.00	50.0%
EarthCube	-	1.50	4.50	3.00	200.0%

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

- CAREER: OCI's CAREER awards support young investigators who exemplify the role of teacher-scholars through outstanding research, excellent education, and the integration of education and research within the context of the mission of their organization. OCI estimates that it will make approximately 5 CAREER awards in FY 2013.

- Cyber-Enabled Materials, Manufacturing, and Smart Systems (CEMMSS): OCI support will fund awards related to computing algorithms and infrastructure, and data systems for the Materials Genome Initiative.

- Cyberinfrastructure Framework for 21st Century Science and Engineering (CIF21): OCI support focuses on the extension of OCI core programs in data, software, networking, and virtual organizations across NSF communities to serve as resources for enhancing scientific collaboration and discovery.

- Comprehensive National Cybersecurity Initiative (CNCI): OCI continues support of CNCI with awards focused on game-changing security prototypes and experimental approaches.

- Expeditions in Education (E^2): Through the E^2 program, OCI will participate in efforts across NSF to assure that scientists and engineers have the training and skills to realize the potential of new cyberinfrastructure tools.

- Innovation Corps (I-Corps): OCI support for I-Corps will fund awards for mentoring researchers or providing additional support for projects that have been previously funded to transition them into practical cyberinfrastructure tools.

- Integrated NSF Support Promoting Interdisciplinary Research and Education (INSPIRE): OCI support of INSPIRE will support Creative Research Awards for Transformative Interdisciplinary Ventures (CREATIV) that help further the OCI goals of creating integrated cyberinfrastructure across

multiple disciplines. OCI will also participate in additional award mechanisms that might be issued as part of the INSPIRE investment.

- Science, Engineering, and Education for Sustainability (SEES): OCI will support computational approaches in climate modeling, system dynamics, and fuel research. OCI supports Research Coordination Networks in the SEES-RCN track to bring different communities together through digital environments to address sustainability problems.

- Secure and Trustworthy Cyberspace (SaTC): OCI support of SaTC is focused on the transition to practice perspective, which leverages successful results from previous and current research and focuses on later stage activities in the research and development lifecycle.

- EarthCube: EarthCube is a partnership between GEO and OCI to create integrated data management infrastructures across the geosciences. OCI funding supports multiple applied areas of cyberinfrastructure development, leveraging OCI core programs, with emphasis on data sharing and archiving, networking, and community building.

Summary and Funding Profile

OCI invests in core research and education, as well as research infrastructure, such as high performance computing resources, software, data, and networking infrastructure. In FY 2013, the number of research grant proposals is expected to increase by 10 from the FY 2012 Estimate as OCI initiates programs to award smaller grants in conjunction with CIF21. The average annualized award size decreases in FY 2013 relative to the FY 2012 Estimate also for this reason. In addition, due to a large acquisition in FY 2012 for the Track 1 HPC program, the average and, to a lesser extent, the median award size in FY 2012 will be larger than in FY 2013.

OCI Funding Profile

	FY 2011 Actual Estimate	FY 2012 Estimate	FY 2013 Estimate
Statistics for Competitive Awards:			
Number of Proposals	706	680	690
Number of New Awards	151	135	145
Funding Rate	21%	20%	21%
Statistics for Research Grants:			
Number of Research Grant Proposals	642	670	680
Number of Research Grants	115	125	135
Funding Rate	18%	19%	20%
Median Annualized Award Size	$127,707	$175,000	$160,000
Average Annualized Award Size	$159,221	$250,000	$225,000
Average Award Duration, in years	2.7	2.5	2.6

Program Monitoring and Assessment

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.

Committees of Visitors (COV):

- In 2011 a COV reviewed OCI. The COV presented its report to the Advisory Committee for Cyberinfrastructure in November 2011. The COV generally found the integrity and efficiency of OCI's processes and management to be high. The COV found the merit review process and portfolio of awards to be excellent. The COV recommended a change to the DataNet program, which is being implemented with a new solicitation in FY 2012. The COV also found that OCI faces unique challenges in creating long-term infrastructure programs with rotating program directors and in developing new communities around data-intensive science while maintaining support for its existing science and engineering computing community.

Number of People Involved in OCI Activities

	FY 2011 Actual Estimate	FY 2012 Estimate	FY 2013 Estimate
Senior Researchers	1,004	775	800
Other Professionals	426	320	330
Postdoctorates	78	85	90
Graduate Students	324	325	320
Undergraduate Students	197	155	155
Total Number of People	**2,029**	**1,660**	**1,695**

Office of Cyberinfrastructure High Performance Computing Portfolio

OCI High Performance Computing Funding

(Dollars in Millions)

	Prior Years[1]	FY 2011 Actual	FY 2012 Estimate[2]	FY 2013 Request
Track 1	329.49	92.13	39.41	30.00
Innovative HPC Program	171.52	4.97	48.90	25.00
Teragrid - Phase III (XD)	23.82	90.81	-	27.00
Total	**$524.83**	**$187.91**	**$88.31**	**$82.00**

Totals may not add due to rounding.

[1] Prior Years includes $17.0 million of ARRA funding in FY 2009 and support for Teragrid (included on the Track 1 funding line) through FY 2010.

[2] FY 2012 Estimate excludes $71.10 million of FY 2011 funds that will be provided for Track 1 in FY 2012. The FY 2011 funds were originally obligated in FY 2011 to fund FY 2013 – FY 2015 of another OCI award (XD) but they were deobligated in FY 2012 in order to complete the Blue Waters funding. No funding is needed for XD in FY 2012 due to forward funding provided in FY 2011.

NSF has been a leader in High Performance Computing (HPC) for almost four decades. Due to the continuing rapid change 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 new coordinated NSF-wide strategy, which is a key component of the 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 to support multidisciplinary computational and data-enabled science and engineering that supports the entire scientific, engineering, and education community. This shift away from a focus on procurement is consistent with the recommendations of a recent review[1] of the Federal Networking and Information Technology R&D (NITRD) program by the President's Council of Advisors on Science and Technology.

TRACK 1 (PETASCALE COMPUTING) – BLUE WATERS

Description

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign has provided the capability for researchers to tackle much larger and more complex research challenges than previously possible, by acquiring, deploying, and operating a sustained petascale leadership class high-performance computational resource known as Blue Waters. This investment complements the Department of Energy (DOE) Office of Science program on computing hardware, which focuses on peak petascale performance, with Blue Waters providing sustained petascale performance. This system, which is located at the University of Illinois at Urbana-Champaign, will be delivered in FY 2012. It is operated by NCSA and its partners in the Great Lakes Consortium for Petascale Computing (GLC).

[1] *Designing a Digital Future: Federally Funded Research and Development Networking and Information Technology,* President's Council of Advisors on Science and Technology, December, 2010.

The Blue Waters project also includes education and outreach programs that will target pre-college, undergraduate, graduate, and post-graduate levels. A Virtual School of Computational Science and Engineering has been 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 will provide 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 will promote 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-oriented and science-oriented careers.

Current Status

In late FY 2010 UIUC was notified by the vendor, IBM, that FY 2011 sub-award milestones would be delayed. UIUC submitted a request to NSF to extend the project end date by approximately nine months, and an external review panel at NSF approved a change in schedule. Under a dispute resolution procedure in the contract, UIUC and IBM attempted to reach agreement regarding project timeline and scope. The parties were unable to reach agreement, and in August 2011, per the terms of their agreement, IBM terminated its contract with UIUC. UIUC submitted a change to the Project Execution Plan (PEP) that was reviewed by an external panel and approved in late September 2011. UIUC selected a new vendor, Cray, for the project, which is expected to be complete by the end of FY 2012.

Science and engineering research and education activities enabled by Blue Waters

This award permits investigators across the country to conduct innovative research demanding petascale capabilities. Allocations have been requested for research on: 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.

Management and Oversight

NSF Structure: The project is managed and overseen by OCI program staff and a grants officer from the Division of Grants and Agreements (DGA). These NSF staff members receive strategic advice from NSF's CIF21 Strategy and Leadership group, which includes representatives from the various 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 oversees work by a number of sub-awardees, conducts software development, and assists competitively selected research groups to prepare to use the Blue Waters system. The primary sub-awardee is responsible for implementation of the hardware, system software, and main program development tools. Other sub-awardees will work on performance modeling, the evaluation of an astrophysical modeling framework, the engagement of applications groups, scalable performance tools, undergraduate training, and broadening the participation of underrepresented groups in high-performance computing. Following system testing and acceptance in 2012, the Blue Waters project will enter a five-year operations phase. A proposal from UIUC for operations is anticipated in FY 2012. The project team is advised by a Petascale Executive Advisory Committee composed of senior personnel with technical and management expertise in high-performance scientific computing, the management of acquisition contracts for leading-edge computing systems, and the operation of large computing centers.

Risks: Any activity of this nature, and at this scale, comes with a certain element of risk. The extensive 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 Track 1 award required 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, 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. Discussion of risks is part of the weekly discussions between UIUC and NSF. Periodic closed session updates to the National Science Board (NSB) identify any major changes in risk assessment.

Reviews: The project was selected through a competitive review in 2007. An external panel of experts, selected by NSF, 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. At the time of writing, these external reviews had been conducted in February 2008, April 2008, October 2008, April 2009, July 2009, December 2009, April 2010, September 2010, December 2010, February 2011, and May 2011, and September 2011.

INNOVATIVE HPC

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 will be made in the context of the eXtreme Digital (XD) services program (described below). While the Track 1 (Petascale Computing) 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 of researchers (over the course of a year) doing leading-edge science and engineering.

There is a direct relationship between the Innovative HPC awards and the XD activity. Several systems are currently serving as allocatable 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 due to 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. Once 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, based on feedback from the scientific and engineering community, a more sustained approach to core HPC services was initiated. This provides a longer time horizon for funding of 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 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.

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; 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 for Innovative HPC

NSF Structure: OCI program officers provide direct oversight during both the acquisition and operations phase. Formal reporting consists of quarterly and annual reports, which are reviewed by the program officer. There are also bi-weekly teleconferences with NSF program officers.

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, by nature, and therefore encompass 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 will be performed as part of the XD review. Semi-annual reviews will be performed as part of the acquisition phase. The reviews are arranged by the NSF program officer. The reviewers' backgrounds include scientific research, project management, and large-scale systems acquisitions and operations, and include familiarity with projects funded by NSF as well as other federal agencies. To the extent possible, continuity through the series of reviews will be provided by using the same set of reviewers.

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 the review process and the NSF program officer.

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

Current Status: Machines and facilities that are currently operational in the Innovative HPC program include Blacklight, Forge, Kraken, Lonestar, Longhorn, Ranger, and Trestles. FutureGrid, Gordon, Keeneland are new highly innovative experimental systems that are becoming fully operational in FY 2012. An award for $27.50 million was made in September 2011 to the University of Texas at Austin for the development of a new system called Stampede. Stampede will be a comprehensive computing, data analysis, and visualization system. It is expected to come online in January 2013 and will be the most powerful system in the NSF XD environment.

The largest resource in the NSF portfolio designed for capacity computing, Kraken, located at the National Institute for Computational Sciences at the University of Tennessee, has reached the end of its useful lifetime. The intent is to provide an additional year of funding to Kraken in FY 2012 while the new solicitation and review process proceeds during late FY 2012 through early FY 2013. New award(s) are expected in FY 2013 to provide the capacity level computational needs of the Open Science community.

FutureGrid Experimental High Performance Grid Testbed at Indiana University (IU)

Description

- The project team, led by Indiana University, has provided a significant new experimental computing grid and cloud test-bed, named FutureGrid, to the research community, together with user support for third-party researchers conducting experiments on FutureGrid. This will enable them to tackle complex research challenges in computer science related to the use and security of grids and clouds.

- The test-bed includes a geographically distributed set of heterogeneous computing systems, a data management system that will hold both metadata and a growing library of software images, and a dedicated network allowing isolatable, secure experiments.

- The test-bed will support virtual machine-based environments, as well as native operating systems for experiments aimed at minimizing overhead and maximizing performance.

- The project partners will integrate existing open-source software packages to create an easy-to-use software environment that supports the instantiation, execution, and recording of grid and cloud computing experiments.

- The FutureGrid system is being integrated into XD in FY 2012 as an operational system.

Gordon Data-Intensive Computing at San Diego Supercomputer Center (SDSC)

Description

- The University of California at San Diego (UCSD) has provided a ground-breaking new computing facility, Gordon, which will be made available to the research community together with advanced user support for researchers with data-intensive problems that may not parallelize well or will require access to very large amounts of memory.

- The distinguishing features are the integration of solid state disks (SSDs) and very large shared memory. This system will be optimized to support research with very large data-sets or very large input-output requirements. It will provide a step-up in capability for data-intensive applications that scale poorly on current large-scale architectures, providing a resource that will enable transformative research in many research domains.

- As a result of the change in the processor timeline, the system will become fully operational as part of the NSF XD environment in FY 2012.

Keeneland Experimental High Performance Computing at Georgia Institute of Technology

Description

- Keeneland, provided by the Georgia Tech Research Corporation (GTRC), is a new experimental high performance computing facility with unconventional computing architectures. The system will allow scientific and engineering researchers to evaluate the relative merit of these new architectures.

- The distinguishing feature of Keeneland is the inclusion of General-Purpose computation on Graphics Processing Units processors (GPGPU) as a general purpose technique for computational acceleration in large systems. The goal is to address computational problems that are challenging to more conventional supercomputing architectures. Productivity is of particular interest in using Open Computing Language (OpenCL) as a mechanism to program the GPGPUs.

- Applications will require additional development and testing to be appropriately prepared to effectively use this new type of architecture.

- An initial system for prototype software and applications development was installed in FY 2010. A full-scale system is being installed in FY 2012 and will become an XD resource.

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

Description

- The Stampede project at the University of Texas at Austin will deliver a new system to NSF XD cyberinfrastructure services in FY 2013, replacing 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 new system will boost XD resources to nearly twice their current capacity. It provides researchers with early access to a potentially transformative new approach to performance via Intel Many Integrated Core (MIC) processors. An addition of the second generation of the MIC processors is planned in late FY 2015.

TERAGRID PHASE III: EXTREME DIGITAL (XD)

Description:

- XD, successor to the TeraGrid program, is an advanced, nationally distributed, open cyberinfrastructure comprised of shared user and management services (XSEDE, eXtreme Science and Engineering Discovery Environment), supercomputing, storage, analysis, and 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.

Science and engineering research and education activities enabled by XD

- XD services will 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, 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 is developing tools and services that not only link users to national facilities but that enable scientific collaboration within and across university campuses, government laboratories, and experimental facilities.

- The project includes outreach and training critical to reducing the barriers to the use of advanced digital systems by the research and education communities. XD 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 inter-related parts – High-Performance Remote Visualization Service (HPRVS); Technology Audit Service (TAS); Technology Insertion Service (TIS); Coordination and Management Service (CMS); Advanced User Support Service (AUSS); and Training, Education and Outreach Service (TEOS). The last three elements constitute the XSEDE project that manages shared services in the XD 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 HPRVS was reviewed in FY 2009 and two awards were made, one to the University of Texas at Austin ($7.0 million) and one to the University of Tennessee-Knoxville ($10.0 million).

- The TAS and TIS components of XD were reviewed in FY 2010 and two awards were made; one award to the University of Buffalo for the TAS and one award to the University of Illinois at Urbana-Champaign for the TIS. These two awards have facilitated the TeraGrid to XD transition and are already functioning successfully in the new environment.

- The final component of XD, called XSEDE, for CMS, AUSS and TEOS, has been awarded to the University of Illinois Urbana-Champaign (UIUC) as of July 1, 2011. This is a five-year award for $121.0 million and involves four other major partners. They are the University of Pittsburgh, University of Texas at Austin, University of California San Diego (UCSD) and the University of Tennessee at Knoxville. XSEDE also involves twelve other institutions.

- Similar to TG, XD will be managed by OCI, informed by the ACCI and its task forces, with ongoing strategic guidance from the NSF cross-directorate CIF21 Strategy and Leadership Group. The project has an external advisory board, a user board, and a service providers board to ensure that all stakeholders can provide project input. These boards provide substantial ongoing community input to the XD project.

- OCI participates in the management of XSEDE via weekly teleconferences with the senior XSEDE personnel.

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.

Current Status: Two planning grants, one to UCSD ($1.60 million) and one to UIUC ($1.62 million), were made in FY 2009 to obtain community input and engagement in order to develop the ideas and expanded horizons that will be required to deploy the advanced infrastructure required for XD. The planning grants were reviewed in February 2010 and the two teams submitted their full proposals in July 2010. The full proposals were reviewed by an external panel of experts in the fourth quarter of FY 2010. A recommendation for an award, entitled XSEDE, was approved by the NSB and awarded July 1, 2011.