OFFICE OF CYBERINFRASTRUCTURE (OCI) $228,070,000
+$13,790,000 / 6.4%

OCI Funding
(Dollars in Millions)

	FY 2009 Omnibus Actual	FY 2009 ARRA Actual	FY 2010 Estimate	FY 2011 Request	Change Over FY 2010 Estimate Amount	Change Over FY 2010 Estimate Percent
Software	$11.96	$41.34	$12.50	$28.77	$16.27	130.2%
Data	11.28	-	12.27	22.04	9.77	79.6%
Other Disciplinary and Interdisciplinary Research	6.44	19.16	23.78	22.31	-1.47	-6.2%
Education	1.35	2.50	10.77	11.21	0.44	4.1%
High Performance Computing	153.93	17.00	113.00	110.30	-2.70	-2.4%
Other Networking and Computational Programs	10.24	-	37.38	28.36	-9.02	-24.1%
Stewardship	4.03	-	4.58	5.08	0.50	10.9%
Total, OCI	**$199.23**	**$80.00**	**$214.28**	**$228.07**	**$13.79**	**6.4%**
Research	29.68	60.50	48.55	73.12	24.57	50.6%
Education	1.35	2.50	10.77	11.21	0.44	4.1%
Infrastructure	164.17	17.00	150.38	138.66	-11.72	-7.8%
Stewardship	4.03	-	4.58	5.08	0.50	10.9%

Totals may not add due to rounding.

OCI supports research, development, acquisition, and operation of advanced shared and connecting cyberinfrastructure (CI) that enables otherwise unrealizable advances in 21st century science and engineering research and education. It increasingly supports the use of advanced CI to address frontier science problems through the growing discipline of computational science and engineering, as well as the computational scientists who develop and use it. OCI capitalizes on a broad range of fundamental scientific and engineering research and education to create and expand the next generation of CI. CI is used in converting data to knowledge, understanding complexity through simulation and prediction, and creating more systematic knowledge about the social and technical issues of large-scale, multidisciplinary collaborative communities, known as virtual organizations, needed to address complex problems and grand challenges facing science and society.

OCI-supported CI includes the comprehensive set of deployable hardware, software, and algorithmic tools that support research, education, and collaboration across and among all research disciplines, whether they are experimental, theoretical, and/or computational. CI consists of computing systems, data storage systems, data repositories, advanced instruments, visualization environments, people, and the necessary intellectual capital, all linked together by software and advanced networks to sustain and improve scholarly productivity and enable breakthroughs in complex problem solving. OCI supports socio-technical research on the use of CI and on ways of improving its effectiveness. It supports training in the development and use of advanced CI as well as research on its use to enhance learning. OCI also supports the scientific and engineering professionals who create and maintain these leading-edge resources and systems, and who provide the Nation's researchers and educators with essential CI services. OCI makes investments that improve CI for science and engineering research, funding the deployment of current CI and innovative developments in future CI. In doing so, it both leverages and complements investments made by other federal agencies. For example, some of NSF's high-end computing investments take advantage of expertise at laboratories funded by the Department of Energy (DOE) and hardware and software developments funded by the Department of Defense's Defense Advanced Research Projects Agency (DARPA). In addition, OCI investments in petascale applications and tools complement those of DOE's Scientific Discovery through Advanced Computing (SciDAC) program, and

OCI's TeraGrid infrastructure is used by researchers funded by the National Institutes of Health (NIH), DOE, and other agencies.

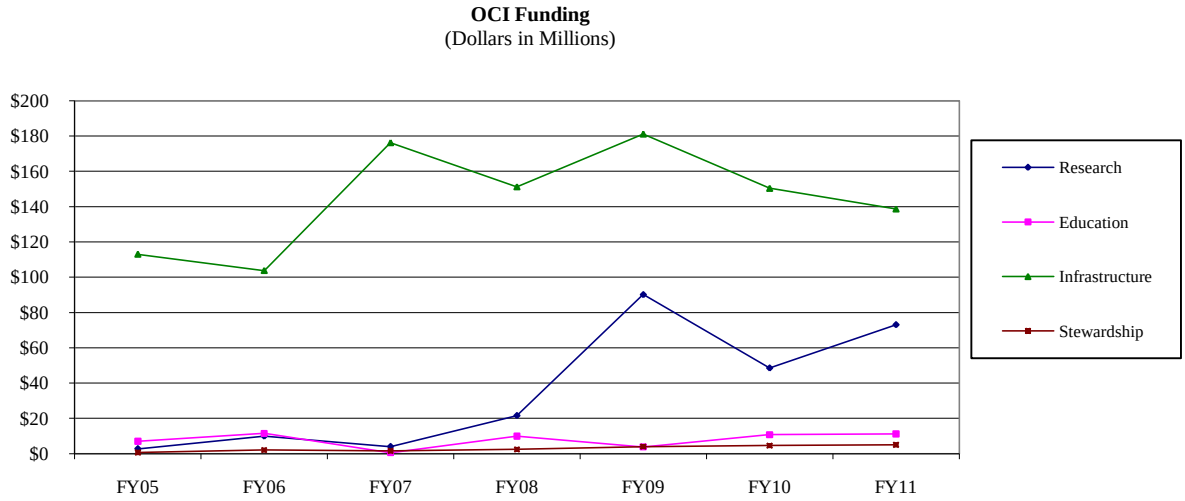

OCI in Context

Science is increasingly data-, compute- and collaboration-driven, requiring access to a diverse set of data, computational resources, storage resources, remote instruments, and services to support their integration and use. Beyond the world-class computational capabilities that already exist, future large-scale and data intensive systems require investment in hardware acquisition, software development and support, and workforce development. Each of these areas requires long term commitments in CI that provide the stability necessary for application support and the incentives for innovation for new capabilities because each area requires multiple years of development in order to achieve the level of performance and maturity useful for science and engineering. Their integration with each other, with discipline specific CI activities, remote instruments, and campus environments supports next generation science activities such as data, visualization, networking, and virtual organizations.

OCI supports the development and deployment of CI that is shared by all scientific and engineering disciplines, making possible potentially transformative basic research in areas such as nanotechnology, physics, chemistry, materials science, sustainable energy, climate/weather, and engineering. It also promotes interoperability between components of CI here in the U.S. and abroad. Approximately two-thirds of NSF's investments in CI are made by the directorates and offices responsible for fundamental domain specific research and education in science and engineering, with the remaining one-third provided by OCI. Through coordinated planning and investments facilitated by NSF's Cyberinfrastructure Council, OCI works to support integrated applications and teams that use advanced CI to solve complex multidisciplinary science and engineering problems, providing economies of scale and scope to ensure that NSF's CI portfolio delivers the highest returns on the Nation's investment.

OCI's investments are guided by NSF's *Cyberinfrastructure Vision for 21st Century Discovery* (www.nsf.gov/dir/index.jsp?org=OCI), a comprehensive CI strategic plan for the Foundation; by many blue-ribbon panel and advisory reports, such as the 2005 President's Information Technology Advisory Committee (PITAC) report (www.nitrd.gov/Pitac/reports/20050609_computational/computational.pdf); by the America COMPETES Act of 2007; and by the opportunities identified by the academic and industrial research community through workshops and white-papers. The America COMPETES Act calls for the Foundation to conduct long-term basic and applied research on high-performance computing and networking. OCI's investments are central to advancing that goal.

OCI activities are key components in the federal government's Networking and Information Technology Research and Development (NITRD) program. The technologies developed and the systems deployed by OCI facilitate discovery and innovation and bolster national competitiveness. PITAC specifically recommended in 2005 that federal agencies reorganize to more effectively support multidisciplinary computational science, which it called the "third pillar" of science and engineering of the 21st century. OCI was created in 2005, and is now specifically taking on a role as steward of computational science activities in coordination with directorates and offices across the Foundation.

Factors Influencing the Allocation Across Major Programs

Research and Education Grants

- The FY 2008 OCI Committee of Visitors (COV) recommended that OCI *"…address needs in other program areas such as networking, data, virtual organizations, learning, and workforce development. Failure to address this imbalance imperils the overall program. OCI has established itself as program with core research strengths, and is developing a workforce to address the challenges faced by the nation*." In response to the COV report, OCI continues to place increased emphasis on data, virtual organizations, learning, and workforce development.
- In FY 2011 OCI will fund data seed grants and other data activities with the goal of addressing challenges in data that have been highlighted in recent reports such as: NSB report 05-40, "Long-Lived Digital Data Collections Enabling Research and Education in the 21st Century," PCAST: "Leadership Under Challenge: Information Technology R&D in a Competitive World," and "Harnessing the Power of Digital Data for Science and Society," report of the Interagency Working Group on Digital Data to the National Science and Technology Council, January 2009.
- Software must be continually refined to reflect, at one end, new advances in the discipline, and at the other end, to adapt and exploit advances in computing hardware. The NSF *Cyberinfrastructure Vision for 21st Century Discovery* outlined the importance of supporting the continued development, expansion, hardening, and maintenance of end-to-end software systems in order to bring the full power of a national cyberinfrastructure to communities of scientists and engineers. An NSF workshop report on *Planning for Cyberinfrastructure Software* recommended that NSF establish new programs to support multidisciplinary projects of substantial size and duration, so as to engage major scientific communities in the application and adoption of CI software within the context of CI goals. A new Software Institutes program, being initiated in FY 2011, will collectively support a vibrant community of partnerships among academia, government laboratories, and industry for the development and stewardship of sustainable software to establish U.S. leadership in the global knowledge economy. This activity will support open science academic communities while the DoE Office of Science (SC) Scientific Discovery through Advanced Computing (SciDAC) Institutes focus primarily on the national laboratories: (see www.scidac.gov/institutes.html).
- In FY 2011 OCI will invest in Cyberinfrastructure Training, Education, Advancement, and Mentoring for Our 21st Century Workforce (CI-TEAM). This program will position the national science and engineering community to engage in integrated research and education activities promoting, leveraging, and utilizing cyberinfrastructure systems, tools, and services. It will prepare current and future generations of scientists, engineers, and educators to design and develop as well as adopt and deploy cyber-based tools and environments for research and learning, both formal and informal. It will expand and enhance the participation in cyberinfrastructure science and engineering activities of diverse groups of people and organizations, with particular emphasis on the inclusion of traditionally underrepresented individuals, institutions (especially minority serving institutions), and communities as both creators and users of cyberinfrastructure.

Networking and Computational Resources and Infrastructure

- To move beyond the current incarnation of cyberinfrastructure requires a new vision for cyberinfrastructure that is significantly more comprehensive and highly integrated so that multiple science communities can more easily work together to address complex problems, enabling collaborative data- and compute-intensive science.
- OCI will increase support for the Innovative High Performance Computing (HPC) program from $10.0 million in FY 2010 to $30.0 million in FY 2011. These awards will include high-risk/high-payoff approaches to HPC systems capable of meeting the growing demands of the open science and engineering communities.
- In FY 2011, funding is also included for eXtreme Digital (XD) – the successor to TeraGrid. Additional information on XD is provided later in this chapter.

The FY 2011 Request for OCI includes $2.0 million to leverage activities aimed at increasing support for transformative research. Examples of potential foci for these investments include the Strategic Technologies for Cyberinfrastructure (STCI) program and EArly-concept Grants for Exploratory Research (EAGER). The primary purpose of the STCI program is to support work leading to the development and/or demonstration of innovative cyberinfrastructure services for science and engineering research and education that fill gaps left by more targeted funding opportunities. In addition, STCI considers highly innovative cyberinfrastructure education, outreach, and training proposals that lie outside the scope of targeted solicitations. The EAGER mechanism supports high-risk, exploratory, and potentially transformative research.

OCI Administration Priority Programs and NSF Investments

OCI Administration Priority Programs and NSF Investments
(Dollars in Millions)

	FY 2009 Omnibus Actual	FY 2009 ARRA Actual	FY 2010 Estimate	FY 2011 Request	Change Over FY 2010 Estimate Amount	Change Over FY 2010 Estimate Percent
Faculty Early Career Development (CAREER)	-	$4.00	$3.71	$3.97	$0.26	7.0%
Graduate Research Fellowships (GRF)	-	2.00	1.00	1.00	-	-
Science, Engineering, and Education for Sustainability (SEES)	N/A	N/A	5.50	5.00	-0.50	-9.1%
Science and Engineering Beyond Moore's Law (SEBML)	-	-	3.00	3.00	-	-

OCI's FY 2011 budget will fund two key NSF programs that support students and early-career researchers. The budget also encourages potentially transformative research and supports critical priorities in global climate change.

Specific OCI investments include:

- The CAREER program remains the primary mechanism for jump-starting junior faculty toward independent careers in research and education. The program has been very successful in supporting traditional areas of science but less so in computational science. In order to address this, OCI will support CAREER awards in computational science across all disciplines. These will support outstanding young faculty who prototype and develop the next generation of cyberinfrastructure and/or apply it to advance their basic science disciplines.

- In response to the Administration's plan to triple the number of NSF's new Graduate Research Fellowship awards by FY 2013, OCI is contributing to NSF's commitment to encourage more highly talented students to pursue graduate education in science and engineering. GRF is widely recognized as a unique fellowship program because it supports the development of world-class scientists and engineers across all science and engineering fields supported by NSF as well as international research activity.

- In FY 2011, OCI will invest $5.0 million in the NSF-wide Science, Engineering, and Education for Sustainability (SEES) portfolio to integrate efforts in climate and energy science and engineering. OCI investments will increase capacity for peer-to-peer collaboration and new modes of education based upon broad and open access to leadership computing; data and information resources; online instruments and observatories; and visualization and collaboration services.

- In conjunction with NSF's other research directorates and offices, OCI will participate in Science and Engineering Beyond Moore's Law (SEBML). SEBML activities include research into new materials, devices, and processes that exploit the capability to create and manipulate specific quantum states and new algorithms that take advantage of hardware and architecture characteristics to deliver maximal total computing power, including those that utilize quantum interactions. Related to both nanotechnology and cyberinfrastructure, it builds on past NSF investments in these areas and energizes them with new directions and challenges.

For more information on Administration priority programs and NSF Investments, please refer to the Overview and NSF-wide Investments sections.

Program Evaluation and Performance Improvement

The Performance Information 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.

The most recent OCI Committee of Visitors met in FY 2008. The COV focused on two specific areas in the context of OCI's four focus areas of High Performance Computing, Data, Virtual Organizations, and Learning and Workforce Development: (1) assessments of the quality and integrity of program operations and program-level technical and managerial matters pertaining to proposal decisions; and (2) comments on how the outputs and outcomes generated by awardees have contributed to the attainment of NSF's mission and strategic outcome goals. The COV made a number of recommendations that OCI is currently addressing.

The Office of Cyberinfrastructure is working with NSF's Advisory Committee for Cyberinfrastructure (ACCI) to gather input from the researchers and educators who use computing and the technologists who develop high-performance computing on future requirements and opportunities for the national CI. The NSF-wide ACCI has established 6 Task Forces and has asked them to address long-term cyberinfrastructure issues. The Task Forces are:

- Campus Bridging;
- Data;
- Grand Challenges;
- High Performance Computing;
- Software and Tools; and
- Work Force Development.

These Task Forces are composed of NSF program officers from each of the NSF research directorates and offices as well as a set of distinguished members from the external science and engineering community. The Task Forces are beginning to explore, discuss, and generate a set of recommendations and ideas.

Number of People Involved in OCI Activities

	FY 2009 Estimate	FY 2009 ARRA Estimate	FY 2010 Estimate	FY 2011 Estimate
Senior Researchers	395	425	455	425
Other Professionals	233	165	205	210
Postdoctorates	19	77	50	55
Graduate Students	111	384	200	210
Undergraduate Students	104	46	75	80
Total Number of People	**862**	**1,097**	**985**	**980**

OCI Funding Profile

	FY 2009 Estimate	FY 2010 Estimate	FY 2011 Estimate
Statistics for Competitive Awards:			
Number of Proposals	337	510	535
Number of New Awards	192	103	107
Regular Appropriation	97	103	107
ARRA	95	-	-
Funding Rate	57%	20%	20%
Statistics for Research Grants:			
Number of Research Grant Proposals	307	490	515
Number of Research Grants	165	97	103
Regular Appropriation	72	97	103
ARRA	93	-	-
Funding Rate	54%	20%	20%
Median Annualized Award Size	$199,743	$225,000	$230,000
Average Annualized Award Size	$568,144	$395,550	$400,000
Average Award Duration, in years	2.7	2.4	2.5

Office of Cyberinfrastructure High Performance Computing Portfolio

OCI High Performance Computing Funding

(Dollars in Millions)

	Prior Years	FY 2008 Actual	FY 2009 Omnibus Actual	FY 2009 ARRA Actual	FY 2010 Estimate	FY 2011 Request
TeraGrid[1]	N/A	$102.93	$30.98	-	-	-
Track 1	20.00	13.65	64.73	-	90.00	48.50
Track 2	75.03	14.19	55.00	-	10.00	6.80
Innovative HPC Program	N/A	N/A	N/A	N/A	10.00	30.00
TeraGrid - Phase III (XD)	-	-	3.22	17.00	3.00	25.00
Total	**$95.03**	**$130.77**	**$153.93**	**$17.00**	**$113.00**	**$110.30**

Totals may not add due to rounding.

[1] Transition from TeraGrid to eXtreme Digital (XD) in FY 2010 - refer to section on XD for more information.

Track 1 – Blue Waters

Description:

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign will provide the capability for researchers to tackle much larger and more complex research challenges than previously possible, across a wide spectrum of domains. NCSA will acquire, deploy, and operate a petascale sustainable, architecturally coherent, innovative, leadership-class, high-performance computational resource, to be known as Blue Waters, for the science and engineering research community. The DoE SC leadership hardware peaks at petascale, while Blue Waters will provide sustained petascale performance. Also, while the DoE microprocessors are commodity processors, the microprocessors in Blue Waters were altered to address the specific needs of the HPC community. This system will be sited at University of Illinois at Urbana-Champaign (UIUC) where it will be operated by NCSA and its partners in the Great Lakes Consortium for Petascale Computing (GLC).

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 will be established to create courses that focus on petascale computing and petascale-enabled science and engineering. The Virtual School will explore new instructional technologies and create courses, curricula, and certificate programs tailored to science and engineering students. It will also sponsor workshops, conferences, summer schools, and seminars.

The project will include an annual series of workshops targeted at the developers of simulation packages and aspiring application developers. In addition, the project will include 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 to accelerate 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.

Science and engineering research and education activities enabled by Blue Waters:
This award will permit 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, 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, 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 advice from NSF's High-Performance Computing Coordinating Group, which includes representatives from the various directorates and offices and is led by OCI. Advice from the Office of General Counsel (OGC) is sought as necessary.

The contract between UIUC and IBM, the principal sub-awardee, includes milestones at which progress by IBM is assessed through a series of deliverables, including software packages and demonstrations, tests of preliminary hardware, simulators, technical specifications, and programmer guides.

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, IBM, is responsible for implementation of the hardware, system software, and main program development tools. Other sub-awardees 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 mid-2011, the Blue Waters project will enter a five-year operations phase. A proposal from UIUC for operations is anticipated in mid-2010. 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, 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, analyzed, and a mitigation plan created and adhered to. 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 bi-weekly discussions between UIUC and NSF.

Reviews: The project was selected through a competitive review in 2007. An external review 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 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 and December 2009 with further reviews planned for April 2010 and at 4-6 month intervals thereafter. In addition, NSF conducts site visits.

Current Status: The project is currently proceeding as planned. External reviews have praised the technical collaboration between the awardee and the primary vendor.

Track 2

The Track 1 system is targeted to provide sustained petascale performance, while the Track 2 systems provide, at most, petascale peak performance. While each Track 2 system is capable of supporting hundreds of researchers (over the course of a year) doing leading-edge science and engineering, the Track 1 system is expected to support on the order of a dozen projects, each capable of producing break-through results as a direct result of having access to such a facility. In FY 2011, funding will be provided for Track 2D awards only. In previous years funding was provided for Track 2A and 2B awards that have transitioned into production TeraGrid resources. No funds will be awarded for Track 2C.

There is a direct relationship between the Track 2 awards and the TeraGrid activity. Track 2 provides the acquisition process for new systems that will become part of TeraGrid. Track 2 awards are made to an institution following an extensive external review process. Proposals submitted consist of two parts: a) an acquisition component (and associated funding) and b) an operations and maintenance component (and associated funding). All award funding goes to the institution. The institution issues a sub-award to the vendor and perhaps other sub-awards as may be deemed necessary. The vendor receives its funding from the NSF awardee following a successful acceptance process agreed upon by NSF. The acquisition phase is completed when there are no further payments due to the vendor. At that point, the Track 2 award transitions into the TeraGrid with the institution taking on the role of a TeraGrid resource provider. Once the institution is integrated into the TeraGrid it has access to the operations and maintenance funding component of the award.

Immediately below is information that is common to the Track 2D program and hence is applicable to all Track 2D awards. Any differences or project-specific information are discussed in that award's section.

Science and engineering research and education activities enabled by Track 2D:

- The complete spectrum of scientific research is supported, including: climate and weather modeling, cosmology and astrophysics, geosciences, physics, chemistry, biology and medicine, earthquake engineering, and mechanical engineering.
- TeraGrid is required to provide evidence of outreach activities that include various education and training opportunities being made available. These are evaluated as part of the annual review process.
- In addition, part of the Track 2D acquisition review process includes an assessment of education and outreach activities being considered.

Management and Oversight for Track 2D:

NSF Structure:

- NSF oversight is provided by OCI program officers who provide direct oversight during both the

acquisition and operations phase and the system integration into the TeraGrid and the follow-on eXtreme Digital (XD) activity.
- Formal reporting consists of quarterly and annual reports. These 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, prior to award, reviews and analyzes the risks as presented in the proposal and identifies any additional risk that should be considered. The Track 2D award process requires that risks be identified, analyzed, and a mitigation plan created and adhered to. 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.

Reviews:
- Annual reviews are performed as part of the TeraGrid annual review.
- Semi-annual reviews are 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 is provided by using the same set of reviewers.

External Structure:
- Each Track 2D award will be managed under a cooperative agreement. Each Track 2D awardee will be responsible for the satisfactory completion of milestones in order for the spending authorization to be raised. Progress will be determined 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:
- The Track 2D cooperative agreement awards were made in FY 2009 and are proceeding as planned.

Track 2Da - Gordon Data Intensive Computing at San Diego Supercomputer Center (SDSC)

Description:
- The University of California at San Diego (UCSD) will provide 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 system will become part of the NSF TeraGrid and the follow-on eXtreme Digital cyberinfrastructure in FY 2011.

Science and engineering research and education activities enabled Gordon:
- It is expected that data-centric/intensive science research will be enabled by this system. 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.
- Examples of scientific challenges that this resource will allow researchers to tackle include the following:

  - De Novo Genome Assembly: Gene sequencers produce information about many small fragments of a genome. Some recent assembly algorithms use a graph-based approach, much more readily executed on a shared-memory system. Using Gordon, researchers will be able to rapidly assemble complex genomes such as mammalian genomes.
  - Astronomy: Modern astronomy databases can be large; for example, the Sloan Digital Sky Survey is approximately six terabytes in size. Typically, the analysis algorithms that researchers use to perform complex searches for astronomical phenomena can be implemented more easily on shared-memory systems. Gordon will enable researchers to load a copy of the Sloan Digital Sky Survey into the flash memory associated with a super-node, greatly extending the types of analyses astronomers can make.
  - Interaction Networks: Interaction networks, graphs representing the relationships between objects, are used in research in areas such as epidemiology, phylogenetics, systems biology, and population biology. These interaction networks can represent relationships between types of data stored in different databases; for example, the combination of social network databases with medical records and genomic profiles to explore questions such as genetic resistance to disease. Gordon will speed analysis of large interaction networks because the databases can be stored on solid-state disks, greatly reducing access time and permitting more complex types of analysis.
- The project will leverage a number of ongoing educational activities at UCSD to expand and diversify the community of users that can utilize this computational resource, including successful outreach programs for women and minorities from underrepresented groups in science and engineering. The project will also create a summer training program for undergraduates.

Track 2Db – Keeneland Experimental High Performance Computing at Georgia Institute of Technology

Description:
- The Georgia Tech Research Corporation (GTRC) will provide a new experimental high performance computing facility with unconventional architectures, Keeneland, to scientific and engineering researchers so they can evaluate the merit of these architectures.
- The distinguishing feature of Keeneland is the inclusion of General-Purpose computation on Graphics Processing Units (GPGPU) processors as general purpose compute accelerators in a sufficiently large system 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 will be deployed that will allow researchers lead time in order to prepare their applications for the full scale system to be installed two years later.
- The system will become part of the NSF TeraGrid cyberinfrastructure in FY 2012.

Science and engineering research and education activities enabled by Keeneland:
- The Georgia Institute of Technology (Georgia Tech) and its partners, the University of Tennessee at Knoxville and the Oak Ridge National Laboratory, will initially acquire and deploy a small, experimental, high-performance computing system. The project team will use this system to develop scientific libraries and programming tools to facilitate the development of science and engineering research applications. The project team will also provide consulting support to researchers who wish to develop applications for the system using OpenCL or to port applications to the system.
- The final system has the potential to support many different science areas. Possible areas of impact include some of the scientific domains in which GPU-based acceleration has already been demonstrated

to have an impact at smaller scale; for example, chemistry and biochemistry, materials science, atmospheric science, and combustion science.
- In addition to providing infrastructure for science and engineering research and education, the project partners will educate and train the next-generation of computational scientists on cutting-edge computing architectures and emerging programming environments using the experimental computing resource.

Track 2Dc - FutureGrid Experimental High Performance Grid Testbed at Indiana University (IU)

Description:
- This project provides researchers with the capability to tackle complex research challenges in computer science related to the use and security of grids and clouds.
- The project team will provide 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.
- 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.
- Part of the FutureGrid evaluation, and part of the review process, will be a determination of its efficacy as a component in XD.

Science and engineering research and education activities enabled by FutureGrid:
- FutureGrid will support research on topics ranging from authentication, authorization, scheduling, virtualization, middleware design, interface design, and cybersecurity, to the optimization of grid-enabled and cloud-enabled computational schemes for researchers in astronomy, chemistry, biology, engineering, atmospheric science, and epidemiology.
- Researchers will be able to measure the overhead of cloud technology by requesting linked experiments on both virtual and bare-metal systems. U.S. scientists will be able to develop and test new approaches to parallel, grid, and cloud computing, and collaborate with international efforts in this area.
- It will provide an experimental platform that accommodates batch, grid, and cloud computing, allowing researchers to attack a range of research questions associated with optimizing, integrating, and scheduling the different service models.
- It will provide a test-bed for middleware development and, because of its private network, allow middleware researchers to do controlled experiments under different network conditions and to test approaches to middleware that include direct interaction with the network control layer.
- It will develop benchmarks appropriate for grid computing, including workflow-based benchmarks derived from applications in astronomy, bioinformatics, seismology, and physics.
- Education and broader outreach activities include the dissemination of curricular materials on the use of FutureGrid, pre-packaged FutureGrid virtual machines configured for particular course modules, and educational modules based on virtual appliance networks and social networking technologies that will focus on education in networking, parallel computing, virtualization, and distributed computing.
- The project will advance education and training in distributed computing at academic institutions with less diverse computational resources. It will do this through the development of instructional resources that include preconfigured environments that provide students with sandboxed virtual clusters. These can be used for teaching courses in parallel, cloud, and grid computing. Such resources will also

provide academic institutions with a simple opportunity to experiment with cloud technology to see if such technology can enhance their campus resources.

Innovative HPC Program

Using lessons learned during the execution of the HPC Track 2 program and informed by the NSF ACCI's High Performance Computing Task Force, a new Innovative HPC program will be created in FY 2010. This program will provide production ready HPC systems as well as opportunities for investigating innovative high-risk/high-payoff approaches to providing the necessary computational resources requested by the science and engineering community. The new program will be aligned with the eXtreme Digital activity, TeraGrid Phase III, in a manner similar to the association between Track 2 and TeraGrid.

TeraGrid Phase III: eXtreme Digital (XD)

Description:

- The TeraGrid (TG), predecessor to XD, is an advanced, nationally distributed, open cyberinfrastructure comprised of 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.
- It enables and supports leading-edge scientific discovery and promotes science and technology education.
- XD takes 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 will enable transformative advances in science and engineering research, in the integration of research and education, and in broadening participation in science and engineering by under-represented groups, by providing researchers and educators with usable access to extreme-scale digital resources, beyond those typically available on most campuses, together with the interfaces, consulting support, and training necessary to facilitate their use.
- XD will provide high-performance computing services, 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 advanced digital systems by the research and education community will be provided. XD will incorporate new ideas and technologies to enable researchers and students to move transparently between local and national resources, substantially lowering the barrier to an effective use of cyberinfrastructure and promoting enhanced productivity.

Management and Oversight:

NSF Structure:

- XD will consist of several inter-related parts: a High Performance Computing and Storage Service (HPCSS), a High-Performance Remote Visualization Service (HPRVS), a Coordination and Management Service (CMS), a Technology Audit and Insertion Service (TAIS), an Advanced User Support Service (AUSS), and a Training, Education and Outreach Service (TEOS).

  - 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 presenting the individual user with a common user environment regardless of where the resources or user are located.
  - The HPCSS, consisting of four to six nodes, will be identified and funded beginning in FY 2011.
  - The HPRVS was reviewed in FY 2009 and two awards were made, one to the University of Texas ($7.0 million) and one to the University of Tennessee ($10.0 million). The TAIS component of XD was reviewed early in FY 2010 and an award is planned for mid FY 2010 in order to impact the existing TeraGrid extension award and to ease the TeraGrid to XD transition.
  - The final phase of XD, involving the other three services, is scheduled to come on line in March 2011. The total anticipated funding for all four services is approximately $25 million in FY 2011.
- Similar to TG, XD will be managed by OCI, informed by the Cyberinfrastructure Council (CIC), the Cyberinfrastructure Coordinating Committee (CICC), and various working groups, in particular the internal high-performance computing working group and an external Science Advisory Board similar to the TG Science Advisory Board.
- OCI will hold weekly teleconferences with XD senior personnel.

External Structure:
- The final configuration of XD will consist of a number of sites, containing 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.

Reviews:
- There is a well planned external review process underway to ensure that the NSF will provide the U.S. scientific and engineering community with the highest quality state of the art computational facilities to carry out world-class science. It is expected that the review process will be completed during FY 2010.

Current Status:
- Phase I – 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 same two organizations have been authorized to submit full proposals subject to further review internally and externally of the deliverables required in the preliminary proposals.
- Phase II – The full proposals will be reviewed by an external panel of experts during FY 2010 and a recommendation will be made. NSF expects to make an award in FY 2011.
- NSF expects to have all components of XD in production between FY 2011 and FY 2013.