Leadership-Class System Acquisition - Creating a Petascale Computing Environment for Science and Engineering

Program Solicitation
NSF 06-573

National Science Foundation

Office of Cyberinfrastructure

Preliminary Proposal Due Date(s) ***(required):***

September 08, 2006

Full Proposal Deadline(s) (due by 5 p.m. proposer's local time):

February 02, 2007

SUMMARY OF PROGRAM REQUIREMENTS

General Information

Program Title:

Leadership-Class System Acquisition - Creating a Petascale Computing Environment for Science and Engineering

Synopsis of Program:

NSF's goal for high performance computing (HPC) in the period 2006-2011 is to enable petascale science and engineering through the deployment and support of a world-class HPC environment comprising the most capable combination of HPC assets available to the academic community. The petascale HPC environment will enable investigations of computationally challenging problems that require computing systems capable of delivering sustained performance approaching 10^{15} floating point operations per second (petaflops) on real applications, that consume large amounts of memory, and/or that work with very large data sets. Among other things, researchers will be able to perform simulations that are intrinsically multi-scale or that involve the simultaneous interaction of multiple processes.

HPC Resource Providers - those organizations willing to acquire, deploy and operate HPC systems in service to the broad science and engineering research and education community - play a key role in the provision and support of a national HPC environment. With this solicitation, NSF requests proposals from organizations, or groups of organizations, willing to serve as a petascale HPC Resource Provider, and who propose to acquire and deploy a new, state-of-the-art, petascale HPC system.

A competitive, petascale HPC system will:

- Enable researchers to work on a range of computationally-challenging science and engineering applications at the frontiers of research;
- Incorporate reliable, robust system software essential to optimal sustained performance;
- Provide a high degree of stability and usability; and,
- Function as a community-driven resource that actively engages the research and education communities in petascale science and engineering.

A robust and effective HPC acquisition process, driven by the requirements of the science and engineering research and education community, is one of the key elements of NSF's HPC strategy. Accordingly, the desired capabilities of the system to be acquired are defined in terms of performance on model problems.

Cognizant Program Officer(s):

- Stephen Meacham, ITR Program Director (On-Detail to OCI), 1145 S, telephone: (703) 292-8970, fax: (703) 292-9060, email: smeacham@nsf.gov
- Jose Munoz, Deputy Office Director/Senior Scientific Advisor, 1145 S, telephone: (703) 292-8970, fax: (703) 292-9060, email: jmunoz@nsf.gov

Applicable Catalog of Federal Domestic Assistance (CFDA) Number(s):

- 47.080 --- Office of Cyberinfrastructure

Award Information

Anticipated Type of Award: Cooperative Agreement

Estimated Number of Awards: 1

Anticipated Funding Amount: $200,000,000 (Up to $200,000,000 over four years, subject to the availability of funds, with $50,000,000 anticipated in FY2007.)

Eligibility Information

Organization Limit:

Proposals may only be submitted by the following:

- Institutions of higher education, non-profit, non-academic organizations directly associated with educational or research activities, and FFRDCs are eligible to submit proposals. For-profit organizations may also participate in proposals, but only as sub-awardees to a submitting organization.

PI Limit:

None Specified

Limit on Number of Proposals per Organization:

1

An eligible organization may submit only one preliminary proposal to this competition. A full proposal may only be submitted by an organization invited to do so after the required preliminary proposals have been reviewed. There is no restriction on the number of preliminary proposals or full proposals in which an organization may appear as a sub-awardee.

Collaborative projects may **only** be submitted as a single preliminary proposal and a single full proposal, in which a single award is being requested. The involvement of partner organizations should be supported through sub-awards administered by the submitting organization.

Limit on Number of Proposals per PI:

None Specified

Proposal Preparation and Submission Instructions

A. Proposal Preparation Instructions

- **Letters of Intent:** Not Applicable
- **Preliminary Proposals:** Submission of Preliminary Proposals is required. Please see the full text of this solicitation for further information.
- **Full Proposal Instructions:** This solicitation contains information that supplements the standard Grant Proposal Guide (GPG) proposal preparation guidelines. Please see the full text of this solicitation for further information

B. Budgetary Information

- **Cost Sharing Requirements:** Cost Sharing is not required by NSF.
- **Indirect Cost (F&A) Limitations:** Not Applicable

- **Other Budgetary Limitations:** Other budgetary limitations apply. Please see the full text of this announcement for further information.

C. Due Dates

- **Preliminary Proposal Due Date(s)** ***(required):***

  September 08, 2006

- **Full Proposal Deadline(s)** (due by 5 p.m. proposer's local time):

  February 02, 2007

Proposal Review Information Criteria

Merit Review Criteria: National Science Board approved criteria. Additional merit review considerations apply. Please see the full text of this announcement for further information.

Award Administration Information

Award Conditions: Additional award conditions apply. Please see the full text of this announcement for further information.

Reporting Requirements: Additional reporting requirements apply. Please see the full text of this announcement for further information.

TABLE OF CONTENTS

I. INTRODUCTION

What are the three-dimensional structures of all of the proteins encoded by an organism's genome and how does structure influence function, both spatially and temporally? What patterns of emergent behavior occur in models of very large societies? How do massive stars explode and produce the heaviest elements in the periodic table? What sort of abrupt transitions can occur in Earth's climate and ecosystem structure? How do these occur and under what circumstances? If we could design catalysts atom-by-atom, could we transform industrial synthesis? What strategies might be developed to optimize management of complex infrastructure systems? What kind of language processing can occur in large assemblages of neurons? Can we enable integrated planning and response to

natural and man-made disasters that prevent or minimize the loss of life and property? These are just some of the important questions that researchers wish to answer using state-of-the-art High-Performance Computing (HPC) systems.

Science and engineering research and education enabled by state-of-the-art HPC tools have a direct bearing on U.S. competitiveness. If investments in HPC are to have long-term impact on research problems of national need, then HPC resources must deliver high performance capability to a wide range of science and engineering applications.

By 2010-2011, it is anticipated that academic researchers will be able to access a rich mix of HPC systems that:

- Deliver sustained performance in the 10 teraflops to 2 petaflops range on a variety of science and engineering codes;
- Are integrated into a national cyberinfrastructure environment; and,
- Are supported at national, regional and/or campus levels.

In this scenario, it is likely that NSF will directly support several systems delivering sustained performance in the 50 to 200 teraflops range across a broad range of science and engineering research applications, and at least one system capable of exceeding one petaflop/s of sustained performance on the most computationally-challenging research codes.

II. PROGRAM DESCRIPTION

Throughout this document, the term "acquisition" and use of the verb "to acquire" refer to the acquisition of hardware and software components of a computing system by an awardee funded in full or in part through an NSF award. **This solicitation is not a request for proposals to provide computing hardware or software directly to NSF.**

The purpose of this solicitation is to generate proposals from organizations, or groups of organizations, that are committed to the acquisition, deployment and subsequent operation of a single, petascale HPC system designed to advance research on the most challenging science and engineering problems. Through a related competition, described in NSF 05-625, mid-range HPC systems, that may be optimized for particular classes of science and engineering research problems, will be funded. Together these competitions will provide a portfolio of HPC resources that is balanced in terms of architecture and capability. Proposals may come from individual organizations or consortia of organizations. In the latter case, one member of a consortium should be designated as the lead institution. The preliminary and full proposals should be submitted by the lead institution. Other organizations should appear as sub-awardees.

A competitive, petascale HPC system will:

- Enable researchers to work on a range of computationally-challenging science and engineering applications at the frontiers of research;
- Incorporate reliable, robust system software essential to optimal sustained performance;
- Provide a high degree of stability and usability; and,
- Function as a community-driven resource that actively engages the research and education communities in petascale science and engineering.

For the purposes of this solicitation, a proposed petascale system may include: *computing hardware*, including processors, caches (if present) and main memory, inter-connects, I/O sub-system(s); *local on-line storage* of sufficient size to support science and engineering research applications that use the full extent of the computing hardware; *archival storage* of a size appropriate to a system of the scale proposed; a wide-area *network connection*; any other hardware typical of a modern supercomputing system; *system software* including, one or more operating systems, one or more file systems, a set of compilers and run-time libraries, software libraries that support access to the full memory model of the system proposed including one that offers an MPI interface, standard operating system and mathematical libraries, debugging and program development tools, system administration and job scheduling software, user accounting software, any other software typical of a modern supercomputing system; either dedicated nodes or small satellite systems that provide for interactive access, code development, job preparation and staging, and system management. Except where otherwise noted, the term “petascale HPC system” will refer to the totality of these components.

In addition to the primary petascale system, a smaller architectural prototype should be deployed prior to the deployment of the full-scale system. Researchers developing applications for the petascale system should be provided with access to this prototype for code development and testing.

The full-scale petascale system must be located at a single geographical site. A "system" distributed across multiple sites will not be considered.

The petascale system deployed as a result of this solicitation will become part of the portfolio of resources supported by NSF for shared use by the broad science and engineering research and education community. Accordingly, the system will complement the capabilities currently provided by existing NSF Resource Provider sites. Allocations for use of the system will be made through the Large Resource Allocation Committee (LRAC) or its successor. (See http://www.teragrid.org/userinfo/guide_access_accounts.html for a description of LRAC.) It is anticipated that the allocation policy will, as at present, emphasize allocating jobs between different HPC resources in a way that tries to match the needs of the job with the capabilities of the resource. It is expected that research challenges that require petascale resources, together with the development and testing of the codes to be used for these, will have highest priority in allocations on the petascale system.

The desired capabilities of the petascale HPC system are described below in terms of the science and engineering research for which it is to be used (see *Research Scope*). To facilitate the preparation of proposals, a set of model problems is described in *Section V., Proposal Preparation and Submission Instructions*. The model problems, or a subset of them, will be used to test the capability of the delivered system and a portion of the funding may be made contingent upon meeting the levels of performance estimated in the associated full proposal (see *Section VII.B. Award Conditions* of this solicitation for more information).

Detailed information on the format to be followed in each proposal submitted in response to this solicitation is provided in *Section V., Proposal Preparation and*

Submission Instructions.

It is anticipated that NSF will receive questions about the solicitation from prospective proposers between the release of the solicitation and the deadline for proposals. Answers to questions that may be of general interest to prospective proposers will be posted on a "Frequently Asked Questions" page accessible through http://www.nsf.gov/div/index.jsp?div=OCI, beginning approximately seven days after the release of the solicitation. Prospective proposers are encouraged to check this page periodically for updates. **All** proposers should check this page within the two weeks prior to the deadline for submitting preliminary proposals and within the two weeks prior to the deadline for submitting full proposals.

Research Scope.

The desired capabilities of the petascale HPC system may be deduced from examples of the research challenges for which it is intended to be used. These include frontier research in the following areas:

- The radiative, dynamic and nuclear physics of stars;
- The physics of supernovae, gamma-ray bursters, binary black-hole systems, and collisions between neutron stars;
- The detailed structure of, and the nature of intermittency in, stratified and unstratified, rotating and non-rotating turbulence in classical and magnetic fluids, and in chemically reacting mixtures;
- The mechanisms of reactions involving large bio-molecules and bio-molecular assemblages, such as enzymes, ribosomes and cellular membranes;
- The bulk properties of matter under extreme conditions by simulating their behavior from first principles;
- Robust, optimal design of engineered systems;
- The nonlinear interactions between cloud systems, weather systems and the Earth's climate;
- The prediction of a protein's three-dimensional structure given only its primary amino acid sequence;
- The dynamics of the Earth's coupled, carbon, nitrogen and hydrologic cycles;
- Strongly correlated systems;
- A molecular understanding of friction and lubrication;
- The internal structure of the Earth from high-resolution, broad-band, global, seismic inversions;
- The design of molecular electronic devices;
- How the magnetic fields of the Earth, gas giants and stars are generated and how they evolve;
- The determination of the potential energy surface of any chemical reaction to an accuracy of 1 kilocalorie/mole and the subsequent dynamical behavior of molecules on that surface;
- The decadal dynamics of the hydrology of large river basins;
- The onset of coronal mass ejections and their interaction with the Earth's magnetic field, including modeling magnetic reconnection and geo-magnetic sub-storms;
- Heterogeneous catalysis on both semiconductor and metal surfaces;
- The formation and evolution of galaxies;
- Understanding the assembly of capsids;
- High-energy physics processes governed by the strong interaction;
- Control of complex systems;
- The properties and instabilities of burning plasmas, and investigations of active magnetic confinement techniques;
- The analysis of very large astronomical datasets;
- The first-principles design of catalysts, pharmaceuticals, and other molecular materials for specificity and efficiency;
- Materials by design;
- Low Mach-number astrophysical flows such as the expulsion of the outer layers of a star to form a planetary nebula;
- The properties of nano-engineered structures;
- The properties of compressible multi-phase flows;
- The coupled dynamics of marine and terrestrial ecosystems and oceanic and atmospheric physics;
- The development of structure in the early cosmos;
- The formation of molecular clouds and pre-stellar cores;
- The interaction of attosecond laser pulse trains with polyatomic molecules; and,
- The interaction between chemical reactivity and fluid dynamics in complex systems such as combustion, atmospheric chemistry, and chemical processing.

It is anticipated that in the 2011-2013 period, the petascale system will be one of only a small number of systems capable of delivering the very large computational power and memory necessary to address these problems. It is important that the design of the system be such that, with perhaps some algorithmic modifications, research codes used in each of these areas be able to achieve a significant fraction of peak performance. Proposers are encouraged to study extant performance measurement results that identify aspects of system design which have been found to be bottlenecks for codes used in the areas of research described above. It is anticipated that the successful proposal will not simply try to maximize the theoretical floating point operation count of the system but will pay careful attention to the flow of data within the extended system, investing resources in enhancing the performance of memory sub-systems, system-wide networking, and I/O.

Budget Scope

Each proposal should be for the acquisition and deployment of a single petascale HPC system for which the four-year project costs total no more than $200,000,000. The award will support the acquisition and deployment of hardware, system software, core software libraries, and personnel costs associated with system acquisition and deployment, including acceptance testing. In addition, in *Section V., Proposal Preparation and Submission Instructions,* proposers are asked to include a plan for how the proposing organization(s) and vendor(s) will work with potential users to prepare their applications to run at large scale on the full system. Up to $3,000,000 of the proposed budget may be used to implement this plan. Detailed budgetary information should be provided in the standard Budget and Budget Justification sections of the preliminary and full proposals as described in *Section V*.

NSF anticipates that, subject to availability of funds, the cost of operating the petascale system, including the cost of providing basic user support, will be covered through one or more awards separate from the award that is the focus of this solicitation. It is anticipated that the award(s) for operational costs will cover the useful life-time of the petascale system after it is fully deployed in production mode. These awards will also cover negotiated operating costs and user support costs that are incurred within the four-year period of the award, if preliminary phases of the system enter production mode prior to deployment of the full system.

The reviewers of preliminary and full proposals will be asked to consider the "total cost of ownership" of the proposed systems. This includes not only the cost of the award, covered by this solicitation, to fund the acquisition and deployment of the petascale system by the awardee, but also the proposing organization's estimate of the costs associated with operating the full system for five years after it enters full production mode and, if a phased deployment is proposed, operating costs associated with preliminary phases. (See *Section V* for information on where to include these estimates.)

III. AWARD INFORMATION

Anticipated Type of Award: Cooperative Agreement.

Estimated Number of Awards: 1

Anticipated Funding Amount: Up to $200,000,000 over four years, subject to the availability of funds, with $50,000,000 anticipated in FY2007.

IV. ELIGIBILITY INFORMATION

Organization Limit:

Proposals may only be submitted by the following:

- Institutions of higher education, non-profit, non-academic organizations directly associated with educational or research activities, and FFRDCs are eligible to submit proposals. For-profit organizations may also participate in proposals, but only as sub-awardees to a submitting organization.

PI Limit:

None Specified

Limit on Number of Proposals per Organization:

1

An eligible organization may submit only one preliminary proposal to this competition. A full proposal may only be submitted by an organization invited to do so after the required preliminary proposals have been reviewed. There is no restriction on the number of preliminary proposals or full proposals in which an organization may appear as a sub-awardee.

Collaborative projects may **only** be submitted as a single preliminary proposal and a single full proposal, in which a single award is being requested. The involvement of partner organizations should be supported through sub-awards administered by the submitting organization.

Limit on Number of Proposals per PI:

None Specified

Additional Eligibility Info:

V. PROPOSAL PREPARATION AND SUBMISSION INSTRUCTIONS

A. Proposal Preparation Instructions

Preliminary Proposals*(required):*

Preliminary proposals submitted in response to this program announcement/solicitation should be prepared and submitted in accordance with the general guidelines contained in the NSF *Grant Proposal Guide* (GPG). The complete text of the GPG is available electronically on the NSF Website at: http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg. Paper copies of the GPG may be obtained from the NSF Publications Clearinghouse, telephone (703) 292-7827 or by e-mail from pubs@nsf.gov.

For **preliminary** proposals, exceptions to GPG guidelines follow.

- The page limit for the Project Description section of the preliminary proposal is ***30 pages.***
- The Budget Justification section is limited to **10 pages**. (Draft copies of sub-contracts or other sub-awards do not need to be included with preliminary proposals.)
- Certain information other than that described in the GPG is to be submitted as Supplementary Documents (see below for details.)
- Collaborative projects may **only** be submitted as a single preliminary proposal, in which a single award is being requested. The involvement of partner organizations should be supported through sub-awards administered by the submitting organization.

Letters of endorsement should not be included with preliminary proposals. Letters of commitment from individuals who are described in the Project Description as involved in the project in a senior capacity but who are not members of the lead proposing organization, or from representatives of institutions or organizations collaborating with the lead institution, are allowable. As described in the Grant Proposal Guide, Section II.C.2.j, such letters of commitment should be included in the Supplementary Documents section and do not count toward overall page limits.

The Project Description must include the following six sections:

- **HPC System Specification**
- **HPC System Performance on Science and Engineering Applications**
- **HPC System Reliability and Usability**
- **Implementation, Project Management, and Risk Mitigation**
- **Quality of the Physical Infrastructure**
- **Effective User Support and Projected Operating Costs**

Information to be provided in each section is described below.

HPC System Specification (Preliminary proposals)

Specify the HPC system to be acquired and deployed by the proposing organization(s). Proposers must provide a detailed description of the technical aspects of the totality of the HPC system proposed, as defined in Section II above, sufficient for reviewers to assess the likely impact on the range of science and engineering research that requires petascale resources. Include a detailed description of any aspects of the proposed system that are likely to influence the performance of science and engineering research codes. Parameters to be considered include total number of processors, speed and architecture of individual processors, number of processors sharing the same access to memory, amount of memory, size and layout of the cache hierarchy, inter-processor and inter-node bandwidth and latency, communications topology, amount of secondary storage, amount of archival storage, I/O sub-system, file system(s), operating system(s), compilers, debugging tools, performance measurement tools, system administration tools. If the system includes any specialized hardware or software to facilitate checkpointing and checkpoint-restart of large jobs, or performance monitoring, please describe this.

Describe how the compute engine, local disk, longer-term mass storage systems, and machine-room network will be integrated to provide an environment that is well configured to support petascale computation.

Describe any vendor-supplied hardware or software support for measuring application and system performance.

HPC System Performance on Science and Engineering Applications (Preliminary proposals)

Provide an explanation of why the proposed system will be suitable for the classes of research problems described in the *Research Scope* sub-section of *Section II, Program Description*. Describe the degree to which the system is designed such that codes of these classes can be expected to scale to large fractions of the total system, and provide relevant justification.

Estimate the time required to boot the full system from a cold start, the time required to start a user's job that requires a large fraction of the available processors, the amount of user memory accessible to a typical processor core without requiring communication over an inter-nodal network, the maximum amount of main memory that will be available to users, and, the time required to exchange the contents of this portion of main memory with local disk storage (both load and store).

HPC System Reliability and Usability (Preliminary proposals)

It is vital that basic system services be sufficient for users and system managers to accomplish their work. System software features of particular importance include the operating system or systems, the file system or systems, compilers, message-passing libraries, other libraries (including standard system and mathematical libraries), debugging tools, application tuning tools, performance monitoring tools, system administration and resource management, job scheduling and accounting, file systems, and software for managing the flow of data within the system.

Describe the software environment that will be delivered with the system and how the proposed system software and tools will respond to the needs of users and system administrators.

Describe the current state of development of system software and tools sufficient to permit effective use of the computational capabilities of the system hardware. Describe the plans for completing this development.

It is anticipated that the job mix for the system will be dominated by jobs that use a significant fraction of the system's resources, accompanied by development jobs that evaluate algorithm performance and scaling for different research problems. It is anticipated that the award document will include a performance requirement similar to the following: that, in production mode, when averaged over one month, 90% of jobs submitted to the system should complete without

having to be resubmitted as a result of a failure in the hardware or system software, including failures as a result of a compiler failing to correctly implement code that complies with the relevant language standard. Include an analysis of the reliability, availability and serviceability properties of the proposed system and the reasons that the proposed system can be expected to meet such a performance requirement.

Implementation, Project Management, and Risk Mitigation (Preliminary proposals)

Provide an implementation plan, including an implementation schedule, and corresponding performance metrics for acquiring and deploying the proposed system. It is anticipated that the start date for the award will be 30th September, 2007. The project implementation plan should include auditable milestones and anticipate progress reviews approximately semi-annually.

A phased deployment plan is acceptable. The implementation plan should be consistent with the full system completing acceptance testing and entering production mode no later than June 30th, 2011. The entry to production mode may be preceded by operations in a friendly-user mode. Describe user access to the system during the deployment phase and prior to system acceptance, including during testing.

To facilitate the rapid use of the new system by researchers, proposers should provide researchers with early access to a prototype, simulator, and tutorials aimed at helping them get applications quickly up and running on the new system.

Staff knowledgeable about the architecture of the petascale system should be available as a resource to prospective users. They should provide expertise on the implications of the system design for application development. Please describe how, in the period prior to the deployment of the full system, the proposing organization(s) and vendor(s) will identify potential users and work with them as they prepare their applications to run at large scale on the full system. Up to $3,000,000 of the acquisition award may be used for this purpose.

Describe the proposed contractual terms of the acquisition sub-contract(s) with the relevant HPC vendor(s). (Copies of draft sub-contracts need not be included with preliminary proposals.)

Describe the availability of experts to address the system integration problems that arise as the system is deployed. This expertise may be provided by the lead proposing organization and/or by sub-awardees. Proposers should make clear their previous associations, if any, with these partners. The breadth of knowledge, depth of interaction, and technical abilities of partners will be considered in the review process. This knowledge and expertise is particularly important in supporting advanced programming paradigms (e.g. compilers for parallel environments, problem solving environments), tools (e.g. performance visualization, parallel debuggers) and system elements (e.g. parallel file systems).

Describe the experience of the proposing organization(s) in the management of large awards and the resources that would be available to manage an award. Describe the experience of the proposing organization(s) in the management of large sub-contracts to vendors for the acquisition of HPC systems. Describe the resources that would be available to manage any such sub-contract issued under an award made as a result of this solicitation. Describe the experience of the primary vendor in the management of hardware and software design and engineering projects for high-performance computing. Describe the project management control mechanisms that will be used by the lead proposing organization and by the primary vendor to manage the project.

Describe the technical and management risks anticipated and the mitigation strategies that will be used.

The Principal Investigator should be the Project Director and be dedicated at least 50% of full-time to the management of the project. It is strongly desired that the Principal Investigator have had formal training and professional experience in the management of large, complex projects. NSF may require that the Principal Investigator complete such training as a condition of an award and/or that the lead proposing organization replace the proposed Principal Investigator with a suitable professional who is given the requisite authority to manage the project. Describe the qualifications and experience of the Principal Investigator with regard to her or his ability to manage a project of this size and complexity, and to manage a resource with a large number of external users. Describe the structure of the management team that will be directly involved in managing the project in its construction and operations phases, the lines of authority, and summarize the relevant qualifications of the individuals making up the management team. (Standard biographical sketches of the people on the management team should be included in the *Biographical Sketches* section of the preliminary proposal.)

Quality of the Physical Infrastructure (Preliminary proposals)

Describe the physical facility that will house the proposed system and any schedule implications of the provision of computer-ready space, including floor space, power, cooling, fire suppression, and any other emergency equipment, for the system and its supporting hardware. Include a description of the physical security that will be provided. Include a description of the expected power and heat budgets of the proposed system and explain how these will be managed. Describe the expected impacts of power interruptions and how these will be managed. Please provide an analysis of the implications of a sudden loss of power to, or catastrophic failure of, either the computing, storage or primary cooling systems and describe what emergency systems will be required to minimize damage to personnel and equipment.

Describe the present external network connectivity and any plans for modifying this prior to system deployment. Describe how the new system will be logically and physically connected to the external network(s).

High-performance applications on the petascale system are expected to produce petabytes of data. Describe how these data will be handled, how data integrity will be maintained, what backup and contingency procedures and schedules will be provided and how will they be implemented. Describe also what issues associated with data access and the protection of confidentiality, privacy and intellectual property are anticipated and how policies and procedures will be developed to deal with these.

Effective User Support and Projected Operating Costs (Preliminary proposals)

Provide a plan for user support that includes a description of the anticipated requirements of the science and engineering research community, a description of

how system resources will be allocated operationally (distinct from the general allocation process handled by LRAC), and any other operational details likely to have an impact on user access or usage of the proposed system. Describe the number and anticipated qualifications of the types of personnel that will be involved with the provision of user support. In addition, describe the user training opportunities that will be made available. Describe the expected availability of dedicated time on the system for both science and engineering applications and systems testing, and the fraction of system resources that will be consumed in moving users on and off the system, or reconfiguring it for dedicated use.

Describe the experience of the proposing organization(s) in operating HPC systems. Include a description of whether operational support was provided on a 24/7 basis or was provided on a more limited basis. Please describe the number and type of users, the types of computation performed, and the nature of the user support provided. Describe the processes used to evaluate management performance, determine user needs, and evaluate user satisfaction.

The reviewers of preliminary proposals will be asked to consider the "total cost of ownership" of the proposed systems. This includes not only the cost of the award, covered by this solicitation, to fund the acquisition and deployment of the petascale system by the awardee but also the proposing organization(s)'s estimate of the costs associated with operating the full system for five years after it enters production mode. In preliminary proposals, proposing organizations should therefore provide, in this section of the Project Description, an analysis of the projected annual operating costs of the proposed system for a period of five years following acceptance of the system, including the cost of providing user support. Detailed operating cost estimates should include: a maintenance contract for 5 years, the annual cost of which does not exceed two per cent of the total hardware acquisition cost; the cost of power and physical security; the cost of network connectivity; costs associated with leasing machine room space, if necessary, and any other necessary operating costs. Provide an estimate of the costs associated with the number of FTEs necessary to maintain 24/7 operations of the proposed system. Provide an estimate of the costs associated with the number of FTEs necessary to provide effective user support. As guidance in preparing these estimates, NSF anticipates that total, fully-burdened, personnel costs for operations, including those involved in providing user support, should not exceed $8,000,000 per year, once the full system is in production mode. If a phased deployment is proposed, estimates of operating costs associated with preliminary phases, after they have entered production mode, should also be included. The estimates of operating costs described in this paragraph should ***not*** be included on the standard FastLane budget forms or discussed in the Budget Justification.

In the preliminary proposals, additional information on operating costs should not be included in the supplementary documents section.

Describe any other factors that are anticipated to have an impact on the Total Cost of Ownership of the proposed system.

Proprietary information (Preliminary proposals)

Proposals containing patentable ideas, trade secrets, privileged or confidential commercial or financial information, disclosure of which may harm the proposer, should be clearly marked where appropriate in the proposal and labeled with the following legend:

“The following is (proprietary or confidential) information that (name of proposing organization) requests not be released to persons outside the Government, except for purposes of review and evaluation.”

Note that proposals submitted to this solicitation will be reviewed by a group of experts that include people who are not U.S. Government personnel.

For further information please refer to the Grant Proposal Guide at http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg

Supplementary Documents (Preliminary proposals)

Preliminary proposals should include the following sections as Supplementary Documents:

- A list of all institutions and companies involved in the project;
- A single, alphabetically ordered list of all people, in the academic or professional computing community, who have collaborated with (within the last 48 months), or have been a Ph.D. advisee or advisor of, any of the personnel involved in the proposed project. In this list, please include, next to the name of each conflicted individual, that individual’s institution or company and the name of the project member with whom he or she has the conflict of interest. It is not necessary to list, as collaborators, personnel who are employees of an institution or company involved in the project; and,
- Letters of commitment from individuals described in the Project Description as involved in the project in a senior capacity, or from representatives of institutions or organizations collaborating with the lead institution.

Proposers are reminded to identify the program announcement/solicitation number (06-573) in the program announcement/solicitation block on the preliminary proposal Cover Sheet. Compliance with this requirement is critical to determining the relevant proposal processing guidelines. Failure to submit this information may delay processing.

Full Proposal Instructions: Proposals submitted in response to this program solicitation should be prepared and submitted in accordance with the general guidelines contained in the NSF Grant Proposal Guide (GPG). The complete text of the GPG is available electronically on the NSF website at: http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg. Paper copies of the GPG may be obtained from the NSF Publications Clearinghouse, telephone (202) 292-7827 or by e-mail from pubs@nsf.gov.

Full proposals may be submitted only by organizations invited to do so after review of the preliminary proposals. The decision to invite or not to invite a full proposal is binding. (See Grant Proposal Guide, Section I.D.2.a.) Full proposals submitted in response to this program announcement/solicitation should be prepared and submitted in accordance with the general guidelines contained in the NSF *Grant Proposal Guide* (GPG). The complete text of the GPG is available electronically on the NSF Website at: http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg. Paper copies of the GPG may be obtained from the NSF Publications Clearinghouse, telephone (703) 292-7827 or by e-mail from pubs@nsf.gov.

For **full** proposals, exceptions to GPG guidelines follow.

- The page limit for the Project Description section of the full proposal is ***60 pages.***
- The Budget Justification section is limited to no more than **60 pages**. It is anticipated that most full proposals will use less than this limit.
- Certain information other than that described in the GPG should be submitted as Supplementary Documents (see below for details.)
- Collaborative efforts may **only** be submitted as a single proposal, in which a single award is being requested. The involvement of partner organizations should be supported through sub-awards administered by the submitting organization.

The Project Description must include the following six sections:

- **HPC System Specification**
- **HPC System Performance on Science and Engineering Applications**
- **HPC System Reliability and Usability**
- **Implementation, Project Management, and Risk Mitigation**
- **Quality of the Physical Infrastructure**
- **Effective User Support and Projected Operating Costs**

Information to be provided in each section is described below.

HPC System Specification (Full proposals)

Specify the HPC system to be acquired and deployed by the proposing organization(s). Proposers must provide a detailed description of the technical aspects of the totality of the HPC system proposed, as defined in Section II above, sufficient for reviewers to assess the likely impact on the range of science and engineering research that requires petascale resources. Include a detailed description of any aspects of the proposed system that are likely to influence the performance of science and engineering research codes. Parameters to be considered include total number of processors, speed and architecture of individual processors, number of processors sharing the same access to memory, amount of memory, size and layout of the cache hierarchy, inter-processor and inter-node bandwidth and latency, communications topology, amount of secondary storage, amount of archival storage, I/O sub-system, file system(s), operating system(s), compilers, debugging tools, performance measurement tools, system administration tools. If the system includes any specialized hardware or software to facilitate checkpointing and checkpoint-restart of large jobs, or performance monitoring, please describe this.

The existence of a clear upgrade path beyond the system proposed to this competition is not a requirement. However, if appropriate, describe how the proposed system is scalable beyond the size proposed for this acquisition or otherwise upgradable in ways that will improve system performance.

Describe how the compute engine, local disk, longer-term mass storage systems, and machine-room network will be integrated to provide an environment that is well configured for petascale computation.

Describe any vendor-supplied hardware or software support for measuring application and system performance.

HPC System Performance on Science and Engineering Applications (Full proposals)

Provide an explanation of why the proposed system will be suitable for the classes of research problems described in the *Research Scope* sub-section of *Section II, Program Description*. Describe the degree to which the system is designed such that codes of these classes can be expected to scale to large fractions of the total system, and provide relevant justification.

In the Project Description, provide a detailed analysis of the anticipated performance of the proposed system on the following set of model problems. Include an estimate of the time to complete each model problem on the system proposed. In the description of each model problem, a target time for completion is included; this is for guidance only.

- A 12288^3 simulation of fully developed homogeneous turbulence in a periodic domain for one eddy turnover time at a value of R_{lambda} of O(2000). The model problem should be solved using a dealiased, pseudospectral algorithm, a fourth-order explicit Runge-Kutta time-stepping scheme, 64-bit floating point (or similar) arithmetic, and a time-step of 0.0001 eddy turnaround times. Full resolution snapshots of the three-dimensional vorticity, velocity and pressure fields should be saved to disk every 0.02 eddy turnaround times. The target wall-clock time for completion is 40 hours.
- A lattice-gauge QCD calculation in which 50 gauge configurations are generated on an 84^3*144 lattice with a lattice

spacing of 0.06 fermi, the strange quark mass m_s set to its physical value, and the light quark mass m_l = 0.05*m_s. The target wall-clock time for this calculation is 30 hours.

- A molecular dynamics (MD) simulation of curvature-inducing protein BAR domains binding to a charged phospholipid vesicle over 10 ns simulation time under periodic boundary conditions. The vesicle, 100 nm in diameter, should consist of a mixture of dioleoylphosphatidylcholine (DOPC) and dioleoylphosphatidylserine (DOPS) at a ratio of 2:1. The entire system should consist of 100,000 lipids and 1000 BAR domains solvated in 30 million water molecules, with NaCl also included at a concentration of 0.15 M, for a total system size of 100 million atoms. All system components should be modeled using the CHARMM27 all-atom empirical force field. The target wall-clock time for completion of the model problem using the NAMD MD package with the velocity Verlet time-stepping algorithm, Langevin dynamics temperature coupling, Nose-Hoover Langevin piston pressure control, the Particle Mesh Ewald algorithm with a tolerance of 1.0e-6 for calculation of electrostatics, a short-range (van der Waals) cut-off of 12 Angstroms, and a time step of 0.002 ps, with 64-bit floating point (or similar) arithmetic, is 25 hours. The positions, velocities, and forces of all the atoms should be saved to disk every 500 timesteps.

In addition, proposers should include, in the Supplementary Documents section of the proposal, an "Application Walk-Through" that describes how applications to run the model problems would be prepared and executed on the proposed system. The description of the Application Walk-Through is included below in the *Supplementary Documents* sub-section.

Proposers should include carefully reasoned predictions of the performance of the full, proposed system on the HPCC benchmark suite.

If reduced-scale prototypes of the proposed system or its component hardware exist, proposers are encouraged to include actual performance data for commonly used benchmarks, particularly the HPCC benchmark suite.

Proposers should include carefully reasoned predictions of the performance of the full, proposed system on the following three non-petascale application benchmarks used in prior solicitations by NSF and other federal agencies. (1) the MILC "Extra-large" problem on 2048 cores, (2) the PARATEC "Large" problem on 256 cores, and (3) the WRF "Large" problem on 512 cores. These are described in NSF-0605, which also provides information on how examples of the benchmark software may be obtained.

Estimate the time required to boot the full system from a cold start, the time required to start a user's job that requires a large fraction of the available processors, the amount of user memory accessible to a typical processor core without requiring communication over an inter-nodal network, the maximum amount of main memory that will be available to users, and, the time required to exchange the contents of this portion of main memory with local disk storage (both load and store).

HPC System Reliability and Usability (Full proposals)

It is vital that basic system services be sufficient for users and system managers to accomplish their work. System software features of particular importance include the operating system or systems, the file system or systems, compilers, message-passing libraries, other libraries (including standard system and mathematical libraries), debugging tools, application tuning tools, performance monitoring tools, system administration and resource management, job scheduling and accounting, file systems, and software for managing the flow of data within the system.

Describe the software environment that will be delivered with the system and how the proposed system software and tools will respond to the needs of users and system administrators. Describe the features of the complete system, in either software or hardware, that will facilitate the development, maintenance and reliable execution of applications that use large fractions of the petascale system.

Please, describe the features of the system software and hardware and, if appropriate, planned operational practices, that will promote fault tolerance. Of particular interest is the impact of fault tolerance on applications that require large fractions of system resources and long run-times.

Describe the current state of development of system software and tools sufficient to permit effective use of the computational capabilities of the system hardware. Describe the plans for completing this development, identifying the personnel and other resources to be allocated to this task and the major accomplishment-based milestones associated with it.

It is anticipated that the job mix for the system will be dominated by jobs that use a significant fraction of the system's resources, accompanied by development jobs that evaluate algorithm performance and scaling for different research problems. It is anticipated that the award document will include a performance requirement similar to the following: that, in production mode, when averaged over one month, 90% of jobs submitted to the system should complete without having to be resubmitted as a result of a failure in the hardware or system software, including failures as a result of a compiler failing to correctly implement code that complies with the relevant language standard. Include an analysis of the reliability, availability and serviceability properties of the proposed system and the reasons that the proposed system can be expected to meet such a performance requirement.

The award instrument will include a performance requirement on the availability of the system. NSF requires that, in production mode, when averaged over a month, the system should be unavailable as a result of scheduled and unscheduled maintenance no more than 8% of the time. Accordingly, provide an analysis of the reasons that the proposed system can be expected to meet this performance requirement.

Implementation, Project Management, and Risk Mitigation (Full proposals)

In the Project Description, provide the following elements of a Project Execution Plan (PEP) for acquiring and deploying the proposed system: a detailed schedule, including auditable, accomplishment-based milestones and corresponding performance metrics; a detailed risk assessment and mitigation plan, identifying both technical and management risks as well as strategies to mitigate such risks; a description of the project management control mechanisms that will be used by the lead proposing organization and by the primary vendor to manage the project. It is anticipated that the start date for the award will be 30th September, 2007. A complete Project Execution Plan should be included in the Supplementary Documents section of the proposal.

A phased deployment plan is acceptable. The implementation plan should be consistent with the full system completing acceptance testing and entering production mode no later than June 30th, 2011. The entry to production mode may be preceded by operations in a friendly user mode. Describe user access to the system during the deployment phase and prior to system acceptance, including during testing.

To facilitate the rapid use of the new system by researchers, proposers should provide researchers with early access to a prototype, simulator, and tutorials aimed at helping them get applications quickly up and running on the new system.

Staff knowledgeable about the architecture of the petascale system should be available as a resource to prospective users. They should provide expertise on the implications of the system design for application development. Please describe how, in the period prior to the deployment of the full system, the proposing organization(s) and vendor(s) will identify potential users and work with them as they prepare their applications to run at large scale on the full system. Up to $3,000,000 of the acquisition award may be used for this purpose

Describe the proposed contractual terms of the acquisition sub-contract(s) with the relevant HPC vendor(s). (As Supplementary Documents, include copies of draft sub-contract(s) from the relevant HPC vendor(s) and any other partners involved that clearly describe the proposed contractual terms of the acquisition and partnership.)

Describe the availability of experts to address the system integration problems that arise as the system is deployed. This expertise may be provided by the lead proposing organization and/or by sub-awardees. Proposers should make clear their previous associations, if any, with these partners. The breadth of knowledge, depth of interaction, and technical abilities of partners will be considered in the review process. This knowledge and expertise is particularly important in supporting advanced programming paradigms (e.g. compilers for parallel environments, problem solving environments), tools (e.g., performance visualization, parallel debuggers) and system elements (e.g., parallel file systems).

Describe the experience of the proposing organization(s) in the management of large awards and the resources that would be available to manage an award. Describe the experience of the proposing organization(s) in the management of large sub-contracts to vendors for the acquisition of HPC systems. Describe the resources that would be available to manage any such sub-contract issued under an award made as a result of this solicitation. Describe the experience of the primary vendor in the management of hardware and software design and engineering projects for high-performance computing.

The Principal Investigator should be the Project Director and be dedicated at least 50% of full-time to the management of the project. It is strongly desired that the Principal Investigator have had formal training and professional experience in the management of large, complex projects. NSF may require that the Principal Investigator complete such training as a condition of an award and/or that the lead proposing organization replace the proposed Principal Investigator with a suitable professional who is given the requisite authority to manage the project. Describe the qualifications and experience of the Principal Investigator with regard to her or his ability to manage a project of this size and complexity, and to manage a resource with a large number of external users. Describe the structure of the management team that will be directly involved in managing the project in its construction and operations phases, the lines of authority, and summarize the relevant qualifications of the individuals making up the management team. (Standard biographical sketches of the people on the management team should be included in the *Biographical Sketches* section of the full proposal.)

Quality of the Physical Infrastructure (Full proposals)

Describe the physical facility that will house the proposed system and any schedule implications of the provision of computer-ready space, including floor space, power, cooling, fire suppression, and any other emergency equipment, for the system and its supporting hardware. Include a description of the physical security that will be provided. Include a description of the expected power and heat budgets of the proposed system and explain how these will be managed. Describe the expected impacts of power interruptions and how these will be managed. Please provide an analysis of the implications of a sudden loss of power to, or catastrophic failure of, either the computing, storage or primary cooling systems and describe what emergency systems will be required to minimize damage to personnel and equipment. Briefly describe the degree to which the physical environment will be compatible with possible system growth through future upgrades. Note that the capability to accommodate future upgrades is not a requirement.

Describe the present external network connectivity and any plans for modifying this prior to system deployment. Describe how the new system will be logically and physically connected to the external network(s).

High-performance applications on the petascale system are expected to produce petabytes of data. Describe how these data will be handled, how data integrity will be maintained, what backup and contingency procedures and schedules will be

provided and how will they be implemented. Describe also what issues associated with data access and the protection of confidentiality, privacy and intellectual property are anticipated and how policies and procedures will be developed to deal with these.

Effective User Support and Projected Operating Costs (Full proposals)

Provide a plan for user support that includes a description of the anticipated petascale resource requirements of the science and engineering research community, a description of how resources will be allocated, and any other operational details likely to have an impact on user access or usage of the proposed system. Describe the number and anticipated qualifications of the types of personnel that will be involved with the provision of user support. In addition, describe the user training opportunities that will be made available. Describe the expected availability of dedicated time on the system for both science and engineering applications and systems testing, and the fraction of system resources that will be consumed in moving users on and off the system, or reconfiguring it for dedicated use.

Describe the experience of the proposing organization(s) in operating HPC systems. Include a description of whether operational support was provided on a 24/7 basis or was provided on a more limited basis. Please describe the number and type of users, the types of computation performed, and the nature of the user support provided. Describe the processes used to evaluate management performance, determine user needs, and evaluate user satisfaction.

Describe the mechanisms that will be used to engage the broad science and engineering community in providing guidance in the operation of the petascale resource and in expanding the scope of petascale research for which it is used.

What activities will be pursued to promote the engagement of traditionally under-represented groups in the use of the petascale system and, more generally, in computational science and engineering?

Describe what metrics will be used to evaluate the effectiveness of operations and user support and to evaluate the degree of success with which the evolving needs of the research community are being identified and met.

The reviewers of full proposals will be asked to consider the "total cost of ownership" of the proposed systems. This includes not only the cost of the award, covered by this solicitation, to fund the acquisition and deployment of the petascale system by the awardee but also the proposing organization(s)'s estimate of the costs associated with operating the full system for five years after it enters production mode. In full proposals, proposing organizations should therefore provide, in this section of the Project Description, an analysis of the projected annual operating costs of the proposed system for a period of five years following acceptance of the system, including the cost of providing user support. Detailed operating cost estimates should include: a maintenance contract for 5 years, the annual cost of which does not exceed two per cent of the total hardware acquisition cost; the cost of power and physical security; the cost of network connectivity; costs associated with leasing machine room space, if necessary, and any other necessary operating costs. Provide an estimate of the costs associated with the number of FTEs necessary to maintain 24/7 operations of the proposed system. Provide an estimate of the costs associated with the number of FTEs necessary to provide effective user support. As guidance in preparing these estimates, NSF anticipates that total, fully-burdened, personnel costs for operations, including those involved in providing user support, should not exceed $8,000,000 per year, once the full system is in production mode. If a phased deployment is proposed, estimates of operating costs associated with preliminary phases, after they have entered production mode, should also be included. ***Information provided will be used to help NSF assess the operating cost-performance attributes of the proposed system.***

The estimates of operating costs described in the preceding paragraph should ***not*** be included on the standard FastLane budget forms or discussed in the Budget Justification. However, if necessary, a more detailed explanation of the estimates of user support and operating costs may be included in the Supplementary Documents section of the proposal.

Describe any other factors that are anticipated to have an impact on the Total Cost of Ownership of the proposed system.

Proprietary information (Full proposals)

Proposals containing patentable ideas, trade secrets, privileged or confidential commercial or financial information, disclosure of which may harm the proposer, should be clearly marked where appropriate in the proposal and labeled with the following legend:

“The following is (proprietary or confidential) information that (name of proposing organization) requests not be released to persons outside the Government, except for purposes of review and evaluation.”

Note that proposals submitted to this solicitation will be reviewed by a group of experts that include people who are not U.S. Government personnel.

For further information please refer to the Grant Proposal Guide at http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg.

Supplementary Documents (Full proposals)

Full proposals should include the following sections as Supplementary Documents:

- Details of the estimated performance on the HPCC benchmark suite, estimated performance on the model problems, and estimated performance on the applications benchmarks described in Section V.A, *System Performance on Science and Engineering Applications (Full proposals)*, of this solicitation. A description of the estimation methodology. This section should not be used to continue discussion or analysis of the merits of the resource provider, vendor or vendors, or system;
- An Application Walk-Through (see below);
- Optionally, additional description of projected operating costs. (***Not to exceed 10 pages, if present;***)
- A Project Execution Plan (PEP) that, at the minimum, should include the following: a summary of the purpose of the facility; a description of the project deliverables; a high-level Work Breakdown Structure (WBS); a Work Breakdown Structure dictionary defining the scope of the WBS elements; a project budget broken out by WBS element; a description of the methodology and assumptions used for estimating the budget components; a project risk analysis and a description of the analysis methodology; a project schedule; a description of the organizational structure of the project team; a description of the subcontracting strategy and controls; a description of the Project Management Control System and financial and business controls to be used; a description of the Configuration Management System to be used; a plan for reporting on the technical and financial status of the project; a description of project governance mechanisms; plans for internal and institutional oversight, external advisory committees, and for building and maintaining effective relationships with the broader research community that will utilize the facility; a description of quality control mechanisms to be used; a description of anticipated safety or health issues associated with the project, if any; plans for systems integration, testing, acceptance, commissioning, and the transition to operational status, including appropriate criteria;
- A list of all institutions and companies involved in the project;
- A single, alphabetically ordered list of all people, in the academic or professional computing community, who have collaborated with (within the last 48 months), or have been a Ph.D. advisee or advisor of, any of the personnel involved in the proposed project. In this list, please include, next to the name of each conflicted individual, that individual's institution or company and the name of the project member with whom he or she has the conflict of interest. It is not necessary to list, as collaborators, personnel who are employees of an institution or company involved in the project;
- Optional letters from science and engineering researchers describing the types of cutting-edge research that they expect to be able to do with the specific system proposed. If included, these letters should be specific about what resources are required and why the proposed system will meet the needs of the problem or problems described. If such letters are included, the page count of all of the letters taken together ***should not exceed 20 pages***. No other types of letters of endorsement are permitted. If other types of letters of endorsement are included or if the page count for endorsement letters is exceeded, NSF may choose to return the proposal without review;
- Optional letters of commitment from individuals described in the Project Description as involved in the project in a senior capacity, or from representatives of institutions or organizations collaborating with the lead institution are allowable; and,
- Copies of draft sub-contracts with vendors, and draft sub-contracts or other sub-award instruments with other partner organizations.

Application Walk-Through (Full proposals)

The Application Walk-Through is intended to be a narrative description of how applications are developed, compiled and run on the system. Three applications should be considered. These correspond to the codes that will be ported or developed to solve the three model problems described above. The goals of the application walk-through are:

- to present compelling evidence that the model problems are thoroughly understood,
- to provide compelling evidence that the model problems will perform on the proposed system as described,
- to demonstrate how the various system components, hardware and software, cooperate to support applications, i.e. show how the components of the petascale system work together as a system.

The application walk-through should be for the petascale system as it is described in the body of the proposal. It should be ***concise*** and should include the following:

1. A complete description of each model problem in sufficient detail to provide context for any claims made in the walk-through. Include any assumptions being made as well as reference material to support any claims and assumptions.

2. An explanation of how the model problems will be implemented on the system.

    - What is the programming paradigm? What languages will be used? What libraries will be required? What does the source code "look like"? (Provide sizing estimates e.g. LOC, bytes.) On what is the source code based?

    - What special features of the system will it use? How significant is this use? How mature and well understood are those features?

    - Provide evidence that the resources and tools to be provided should be expected to achieve the performance being predicted.

- How will the applications that implement the model problems be compiled and linked? What is the programming environment? What libraries are required and how were they generated/developed? Are they commonly accepted and recognized libraries or something new?

3. Describe the parts of the full system that are required for each model problem. If an application will not exercise all components of the system, describe what fraction of each sub-system will be used.

4. Describe, for the full system, a system cold start from power-off to fully functional and ready to launch jobs. What is the POST (Power-On Self Test) sequence? How does the RAS database (if there is one) obtain configuration information? How are failed components recognized and dealt with and how does the system recover and under what conditions does it not recover? What is the OS boot process? What system services are required to start after OS boot and how does the system configuration database get modified to reflect the status of these services? How long will a complete cold-boot process take? How long will a warm boot process take, if one is provided?

5. How will jobs be submitted to the system? What authentication/security mechanisms are provided?

    - How is an interactive parallel job initiated (what is the command line) and what is the step-by-step process of how is it launched? How is the binary loaded into the memory and threads/processes initiated? From what storage/file system? How are resources required by an application identified and made known to the system?

    - How would the system determine exactly which computational units, memory sub-system, IO resources are to be allocated to each of the model problem applications?

    - How are these processes different for launching a batch job?

6. What is the parallel I/O model for your system and how does it work?

    - How does the file system work? What would be the steps for parallel input and output to/from the applications that implement each of the model problems? What software components are required for communication in the IO path?

    - How does a parallel application handle: STDIN/STDOUT/STDERR? What are the steps for data written to STDOUT to be presented to the user during an interactive job running from a shell?

    - How does the user terminate a large parallel application and what are the steps involved in sending and servicing the signal?

    - What mechanisms are provided for application-to-application communication if needed?

7. Addressing execution abnormalities:

    - What is the approach to checkpoint/restart? How much time would it be expected to take and what resources will it utilize?

    - A thread generates a NAN or some other error. How does the system respond and reliably terminate the application?

    - A system node just experienced an unrecoverable memory error. How does the system respond? How does the application get reliably terminated?

    - A link goes down. How does the runtime detect this and respond?

    - How are other aborts handled?

    - How is the system administrator informed of any abnormalities?

8. Performance tuning: Assuming correct results are obtained but performance of the application appears to be lower than expected based on what a simple calculation of bandwidth and operations would indicate or from previous scaling runs, how does the user determine what went wrong and how to recover performance (if possible)?

    - What tools are used and what hardware resources (e.g. hardware performance counters) can be utilized to track down the problems?

    - How are applications debugged? What mechanisms will be provided to assist the user in debugging an application?

What mechanisms will be provided to assist the user in application scaling?

Additional Single Copy Documents

A document describing physical security and cybersecurity plans should be provided. This will be considered part of the PEP but should be uploaded as a separate PDF file under "Additional Single Copy Documents".

Proposers are reminded to identify the program solicitation number (Populated with NSF Number at Clearance) in the program solicitation block on the NSF Cover Sheet For Proposal to the National Science Foundation. Compliance with this requirement is critical to determining the relevant proposal processing guidelines. Failure to submit this information may delay processing.

B. Budgetary Information

Cost Sharing: Cost sharing is not required by NSF in proposals submitted under this Program Solicitation.

Other Budgetary Limitations:

Each proposal should be for the acquisition and deployment of a single, petascale HPC system at a single geographical site for which the four-year project costs total no more than $200,000,000. Each award will support the acquisition and deployment, by the awardee, of hardware, system software, and core software libraries, and personnel costs associated with system acquisition and deployment, including acceptance testing. In addition, up to $3,000,000 of the overall budget may be allocated to activities in which, prior to full-scale system deployment, the proposing organization(s) and vendors(s) identify potential users and work with them to prepare their applications to run at large scale on the petascale system. Detailed budgetary information should be provided in the Budget Justification section of the proposal.

C. Due Dates

- **Preliminary Proposal Due Date(s)** ***(required):***

  September 08, 2006

- **Full Proposal Deadline(s)** (due by 5 p.m. proposer's local time):

  February 02, 2007

D. FastLane Requirements

Proposers are required to prepare and submit all proposals for this program solicitation through the FastLane system. Detailed instructions for proposal preparation and submission via FastLane are available at: http://www.fastlane.nsf.gov/a1/newstan.htm. For FastLane user support, call the FastLane Help Desk at 1-800-673-6188 or e-mail fastlane@nsf.gov. The FastLane Help Desk answers general technical questions related to the use of the FastLane system. Specific questions related to this program solicitation should be referred to the NSF program staff contact(s) listed in Section VIII of this solicitation.

Submission of Electronically Signed Cover Sheets. The Authorized Organizational Representative (AOR) must electronically sign the proposal Cover Sheet to submit the required proposal certifications (see Chapter II, Section C of the Grant Proposal Guide for a listing of the certifications). The AOR must provide the required electronic certifications within five working days following the electronic submission of the proposal. Proposers are no longer required to provide a paper copy of the signed Proposal Cover Sheet to NSF. Further instructions regarding this process are available on the FastLane Website at: http://www.fastlane.nsf.gov

VI. NSF PROPOSAL PROCESSING AND REVIEW PROCEDURES

Proposals received by NSF are assigned to the appropriate NSF program and, if they meet NSF proposal preparation requirements, for review. All proposals are carefully reviewed by a scientist, engineer, or educator serving as an NSF Program Officer, and usually by three to ten other persons outside NSF who are experts in the particular fields represented by the proposal. These reviewers are selected by Program Officers charged with the oversight of the review process. Proposers are invited to suggest names of persons they believe are especially well qualified to review the proposal and/or persons they would prefer not review the proposal. These suggestions may serve as one source in the reviewer selection process at the Program Officer's discretion. Submission of such names, however, is optional. Care is taken to ensure that reviewers have no conflicts with the proposer.

A. NSF Merit Review Criteria

All NSF proposals are evaluated through use of the two National Science Board (NSB)-approved merit review criteria: intellectual merit and the broader impacts of

the proposed effort. In some instances, however, NSF will employ additional criteria as required to highlight the specific objectives of certain programs and activities.

The two NSB-approved merit review criteria are listed below. The criteria include considerations that help define them. These considerations are suggestions and not all will apply to any given proposal. While proposers must address both merit review criteria, reviewers will be asked to address only those considerations that are relevant to the proposal being considered and for which the reviewer is qualified to make judgements.

What is the intellectual merit of the proposed activity?
How important is the proposed activity to advancing knowledge and understanding within its own field or across different fields? How well qualified is the proposer (individual or team) to conduct the project? (If appropriate, the reviewer will comment on the quality of the prior work.) To what extent does the proposed activity suggest and explore creative and original concepts? How well conceived and organized is the proposed activity? Is there sufficient access to resources?

What are the broader impacts of the proposed activity?
How well does the activity advance discovery and understanding while promoting teaching, training, and learning? How well does the proposed activity broaden the participation of underrepresented groups (e.g., gender, ethnicity, disability, geographic, etc.)? To what extent will it enhance the infrastructure for research and education, such as facilities, instrumentation, networks, and partnerships? Will the results be disseminated broadly to enhance scientific and technological understanding? What may be the benefits of the proposed activity to society?

NSF staff will give careful consideration to the following in making funding decisions:

Integration of Research and Education
One of the principal strategies in support of NSF's goals is to foster integration of research and education through the programs, projects, and activities it supports at academic and research institutions. These institutions provide abundant opportunities where individuals may concurrently assume responsibilities as researchers, educators, and students and where all can engage in joint efforts that infuse education with the excitement of discovery and enrich research through the diversity of learning perspectives.

Integrating Diversity into NSF Programs, Projects, and Activities
Broadening opportunities and enabling the participation of all citizens -- women and men, underrepresented minorities, and persons with disabilities -- is essential to the health and vitality of science and engineering. NSF is committed to this principle of diversity and deems it central to the programs, projects, and activities it considers and supports.

Additional Review Criteria:

Proposals for this solicitation will also be subject to the additional review criteria described below. These criteria parallel specific sections in the Project Description.

- **Meeting the Needs of the Most Computationally-Challenging Science and Engineering Applications.** How well does the system proposed match the requirements of the science and engineering research and education community for petascale HPC resources and services? Can the proposed system provide the necessary leadership-class computational capability required to generate new, breakthrough, science and engineering discoveries? In addition, for full proposals, are the estimates of performance on the model problems credible?
- **System Reliability and Usability**. Are the system and the operation of the system likely to provide a robust, reliable, high-productivity computational environment for users? Does the environment in which the system will be embedded include adequate capability for the remote analysis of output from high-end computations? In addition, for full proposals, assess the commitment of the vendor or vendors to meet the agreed performance goals and to provide post-acquisition support.
- **Implementation, Project Management and Risk Mitigation**. Is there an adequate procedure for ensuring that the proposed system will be available for use by the science and engineering research and education community? Does the proposing organization have the capability to manage the award and any associated sub-awards? Does the PI have the capability to manage the project? Is the implementation plan for acquisition and deployment adequate and realistic? Does the proposing organization and its partners have the expertise to meet any challenges likely to be encountered while deploying the complete system (including data storage, communications and core software environment) and bringing it to production status? Has there been a reasonable assessment of potential risks and does the proposal include adequate risk management strategies? Is the Project Execution Plan well designed and will it be effective? .
- **Quality and Availability of the Physical Infrastructure**. Are the physical facilities described by the proposing organization adequate to accommodate the petascale system proposed? Is the physical infrastructure sufficient to mitigate the risk of potential hazards?
- **Effective User Support.** What are the qualifications and experience of the PI and the proposing organization in regard to managing a resource for national use and providing effective user support? Is the user support plan well designed for the support of users with petascale applications? Does the proposal identify appropriate personnel for user support (either currently on staff or to be hired)?
- **Total Cost of Ownership.** Are the budget and roster of personnel for operations and user support adequate and reasonable? Assess the total cost of ownership of the proposed HPC system. Is this reasonable in light of the advances in science and engineering that are likely to result?

B. Review and Selection Process

Proposals submitted in response to this program solicitation will be reviewed by Adhoc Review or Panel Review or Site Visit Review.

Review process and time-line

Preliminary proposals will be reviewed in Fall 2006. NSF will subsequently invite some or all of the proposing organizations to submit full proposals. It is

anticipated that organizations that submitted preliminary proposals will be notified whether they are or are not invited to submit full proposals approximately two months after the deadline for submission of preliminary proposals. Only those organizations invited to submit full proposals will be eligible to do so. NSF expects that the process of reviewing full proposals will include site visits to one or more relevant sites associated with one or more of the full proposals. It is anticipated that these site visits will take place in or around April, 2007. NSF may request that vendors participate in the site visits. During the review process, NSF may submit requests for clarification of aspects of the proposals to one or more of the submitting organizations.

Reviewers will be asked to formulate a recommendation to either support or decline each proposal. The Program Officer assigned to manage the proposal's review will consider the advice of reviewers and will formulate a recommendation.

After scientific, technical and programmatic review and consideration of appropriate factors, the NSF Program Officer recommends to the cognizant Division Director whether the proposal should be declined or recommended for award. NSF is striving to be able to tell applicants whether their proposals have been declined or recommended for funding within six months. The time interval begins on the date of receipt. The interval ends when the Division Director accepts the Program Officer's recommendation.

A summary rating and accompanying narrative will be completed and submitted by each reviewer. In all cases, reviews are treated as confidential documents. Verbatim copies of reviews, excluding the names of the reviewers, are sent to the Principal Investigator/Project Director by the Program Officer. In addition, the proposer will receive an explanation of the decision to award or decline funding.

In all cases, after programmatic approval has been obtained, the proposals recommended for funding will be forwarded to the Division of Grants and Agreements for review of business, financial, and policy implications and the processing and issuance of a grant or other agreement. Proposers are cautioned that only a Grants and Agreements Officer may make commitments, obligations or awards on behalf of NSF or authorize the expenditure of funds. No commitment on the part of NSF should be inferred from technical or budgetary discussions with a NSF Program Officer. A Principal Investigator or organization that makes financial or personnel commitments in the absence of a grant or cooperative agreement signed by the NSF Grants and Agreements Officer does so at their own risk.

VII. AWARD ADMINISTRATION INFORMATION

A. Notification of the Award

Notification of the award is made to *the submitting organization* by a Grants Officer in the Division of Grants and Agreements. Organizations whose proposals are declined will be advised as promptly as possible by the cognizant NSF Program Division administering the program. Verbatim copies of reviews, not including the identity of the reviewer, will be provided automatically to the Principal Investigator. (See section VI.A. for additional information on the review process.)

B. Award Conditions

An NSF award consists of: (1) the award letter, which includes any special provisions applicable to the award and any numbered amendments thereto; (2) the budget, which indicates the amounts, by categories of expense, on which NSF has based its support (or otherwise communicates any specific approvals or disapprovals of proposed expenditures); (3) the proposal referenced in the award letter; (4) the applicable award conditions, such as Grant General Conditions (NSF-GC-1); * or Federal Demonstration Partnership (FDP) Terms and Conditions * and (5) any announcement or other NSF issuance that may be incorporated by reference in the award letter. Cooperative agreement awards also are administered in accordance with NSF Cooperative Agreement Terms and Conditions (CA-1). Electronic mail notification is the preferred way to transmit NSF awards to organizations that have electronic mail capabilities and have requested such notification from the Division of Grants and Agreements.

*These documents may be accessed electronically on NSF's Website at http://www.nsf.gov/home/grants/grants_gac.htm. Paper copies may be obtained from the NSF Publications Clearinghouse, telephone (703) 292-7827 or by e-mail from pubs@nsf.gov.

More comprehensive information on NSF Award Conditions is contained in the NSF *Grant Policy Manual* (GPM) Chapter II, available electronically on the NSF Website at http://www.nsf.gov/cgi-bin/getpub?gpm. The GPM is also for sale through the Superintendent of Documents, Government Printing Office (GPO), Washington, DC 20402. The telephone number at GPO for subscription information is (202) 512-1800. The GPM may be ordered through the GPO Website at http://www.gpo.gov.

Special Award Conditions:

Additional award conditions appropriate for the proper management and oversight of a complex agreement will be negotiated with the proposing organization prior to award and incorporated into the special terms and conditions of the award. It is anticipated that these will be dependent on the nature of the project proposed and the structure of the Project Execution Plan.

C. Reporting Requirements

For all multi-year grants (including both standard and continuing grants), the Principal Investigator must submit an annual project report to the cognizant Program Officer at least 90 days before the end of the current budget period. (Some programs or awards require more frequent project reports). Within 90 days after expiration of a grant, the PI also is required to submit a final project report.

Failure to provide the required annual or final project reports will delay NSF review and processing of any future funding increments as well as any pending

proposals for that PI. PIs should examine the formats of the required reports in advance to assure availability of required data.

PIs are required to use NSF's electronic project-reporting system, available through FastLane, for preparation and submission of annual and final project reports. Such reports provide information on activities and findings, project participants (individual and organizational) publications; and, other specific products and contributions. PIs will not be required to re-enter information previously provided, either with a proposal or in earlier updates using the electronic system. Submission of the report via FastLane constitutes certification by the PI that the contents of the report are accurate and complete.

Additional reporting requirements will be negotiated with the proposing organization prior to award and incorporated into the special terms and conditions of the award. At a minimum, these are likely to include, monthly reports that include reports from the Project Management Control System, more detailed quarterly or semi-annual reports, site visits, and periodic project reviews. It is anticipated that the latter will be approximately semi-annually but the review schedule will be based on the negotiated project schedule.

VIII. AGENCY CONTACTS

General inquiries regarding this program should be made to:

- Stephen Meacham, ITR Program Director (On-Detail to OCI), 1145 S, telephone: (703) 292-8970, fax: (703) 292-9060, email: smeacham@nsf.gov
- Jose Munoz, Deputy Office Director/Senior Scientific Advisor, 1145 S, telephone: (703) 292-8970, fax: (703) 292-9060, email: jmunoz@nsf.gov

For questions related to the use of FastLane, contact:

- FastLane Help Desk, telephone: 1-800-673-6188; e-mail: fastlane@nsf.gov.
- Priscilla L Bezdek, Program and Technology Specialist, 1145 S, telephone: (703) 292-8962, fax: (703) 292-9060, email: pbezdek@nsf.gov

IX. OTHER INFORMATION

The NSF Website provides the most comprehensive source of information on NSF Directorates (including contact information), programs and funding opportunities. Use of this Website by potential proposers is strongly encouraged. In addition, MyNSF (formerly the Custom News Service)is an information-delivery system designed to keep potential proposers and other interested parties apprised of new NSF funding opportunities and publications, important changes in proposal and award policies and procedures, and upcoming NSF Regional Grants Conferences. Subscribers are informed through e-mail or the user's Web browser each time new publications are issued that match their identified interests. MyNSF also is available on NSF's Website at http://www.nsf.gov/mynsf/.

Grants.gov provides an additional electronic capability to search for Federal government-wide grant opportunities. NSF funding opportunities may be accessed via this new mechanism. Further information on Grants.gov may be obtained at http://www.grants.gov.

Answers to questions that may be of general interest to prospective proposers will be posted on a “Frequently Asked Questions” page accessible through http://www.nsf.gov/div/index.jsp?div=OCI, beginning approximately seven days after the release of the solicitation. Prospective proposers are encouraged to check this page periodically for updates. **All** proposers should check this page within the two weeks prior to the deadline for submitting preliminary proposals and within the two weeks prior to the deadline for submitting full proposals.

ABOUT THE NATIONAL SCIENCE FOUNDATION

The National Science Foundation (NSF) is an independent Federal agency created by the National Science Foundation Act of 1950, as amended (42 USC 1861-75). The Act states the purpose of the NSF is "to promote the progress of science; [and] to advance the national health, prosperity, and welfare by supporting research and education in all fields of science and engineering."

NSF funds research and education in most fields of science and engineering. It does this through grants and cooperative agreements to more than 2,000 colleges, universities, K-12 school systems, businesses, informal science organizations and other research organizations throughout the US. The Foundation accounts for about one-fourth of Federal support to academic institutions for basic research.

NSF receives approximately 40,000 proposals each year for research, education and training projects, of which approximately 11,000 are funded. In addition, the Foundation receives several thousand applications for graduate and postdoctoral fellowships. The agency operates no laboratories itself but does support National Research Centers, user facilities, certain oceanographic vessels and Antarctic research stations. The Foundation also supports cooperative research between universities and industry, US participation in international scientific and engineering efforts, and educational activities at every academic level.

Facilitation Awards for Scientists and Engineers with Disabilities provide funding for special assistance or equipment to enable persons with disabilities to work on NSF-supported projects. See Grant Proposal Guide Chapter II, Section D.2 for instructions regarding preparation of these types of proposals.

The National Science Foundation has Telephonic Device for the Deaf (TDD) and Federal Information Relay Service (FIRS) capabilities that enable individuals with hearing impairments to communicate with the Foundation about NSF programs, employment or general information. TDD may be accessed at (703) 292-5090 and (800) 281-8749, FIRS at (800) 877-8339.

The National Science Foundation Information Center may be reached at (703) 292-5111.

The National Science Foundation promotes and advances scientific progress in the United States by competitively awarding grants and cooperative agreements for research and education in the sciences, mathematics, and engineering.

To get the latest information about program deadlines, to download copies of NSF publications, and to access abstracts of awards, visit the NSF Website at http://www.nsf.gov

- **Location:** 4201 Wilson Blvd. Arlington, VA 22230
- **For General Information** (NSF Information Center)**:** (703) 292-5111
- **TDD (for the hearing-impaired):** (703) 292-5090
- **To Order Publications or Forms:**

  Send an e-mail to: pubs@nsf.gov

  or telephone: (703) 292-7827
- **To Locate NSF Employees:** (703) 292-5111

PRIVACY ACT AND PUBLIC BURDEN STATEMENTS

The information requested on proposal forms and project reports is solicited under the authority of the National Science Foundation Act of 1950, as amended. The information on proposal forms will be used in connection with the selection of qualified proposals; and project reports submitted by awardees will be used for program evaluation and reporting within the Executive Branch and to Congress. The information requested may be disclosed to qualified reviewers and staff assistants as part of the proposal review process; to proposer institutions/grantees to provide or obtain data regarding the proposal review process, award decisions, or the administration of awards; to government contractors, experts, volunteers and researchers and educators as necessary to complete assigned work; to other government agencies or other entities needing information regarding applicants or nominees as part of a joint application review process, or in order to coordinate programs or policy; and to another Federal agency, court, or party in a court or Federal administrative proceeding if the government is a party. Information about Principal Investigators may be added to the Reviewer file and used to select potential candidates to serve as peer reviewers or advisory committee members. See Systems of Records, NSF-50, "Principal Investigator/Proposal File and Associated Records," 69 Federal Register 26410 (May 12, 2004), and NSF-51, "Reviewer/Proposal File and Associated Records, " 69 Federal Register 26410 (May 12, 2004). Submission of the information is voluntary. Failure to provide full and complete information, however, may reduce the possibility of receiving an award.

An agency may not conduct or sponsor, and a person is not required to respond to, an information collection unless it displays a valid Office of Management and Budget (OMB) control number. The OMB control number for this collection is 3145-0058. Public reporting burden for this collection of information is estimated to average 120 hours per response, including the time for reviewing instructions. Send comments regarding the burden estimate and any other aspect of this collection of information, including suggestions for reducing this burden, to:

Suzanne H. Plimpton
Reports Clearance Officer
Division of Administrative Services
National Science Foundation
Arlington, VA 22230