

Distributed Terascale Facility (DTF)

Program Solicitation

NSF 01-51

DIRECTORATE FOR COMPUTER AND INFORMATION SCIENCE AND ENGINEERING
DIVISION OF ADVANCED COMPUTATIONAL INFRASTRUCTURE AND RESEARCH

FULL PROPOSAL DEADLINE(S): April 19, 2001



NATIONAL SCIENCE FOUNDATION



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 Web Site 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 or (800) 281-8749
- **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

SUMMARY OF PROGRAM REQUIREMENTS

GENERAL INFORMATION

Program Title: Distributed Terascale Facility (DTF)

Synopsis of Program: In FY 2001 NSF seeks to open a pathway to future computing, communications, and information environments by creating a very large-scale system that is part of the rapidly expanding computational Grid¹. NSF will establish an advanced, multi-site "distributed facility" connected by ultra high-speed networking that will lead to breakthroughs and enhance the capabilities of U.S. researchers in all areas of computational, computer, and information science and engineering. This environment will include at least one single-site computing system capable of five or more teraflops per second (peak) performance. However, since modern scientific and engineering research requires more than just computational capability, this terascale computer system will be embedded within an overall system that also provides sophisticated data handling and interaction with remote sites. This distributed facility will include substantial support for accessing, analyzing, processing, transmitting, and visualizing multi-terabyte data collections of current and future interest to the U.S. research community. This will require the DTF to have terabytes to petabytes of online and archival storage available for user access and multi-gigabit per second network connectivity. The DTF will be fully coordinated with the resources and activities of the existing PACI partnerships. Special consideration will be given to qualified proposals that utilize newer generation processors and other High Performance Computing equipment. Full exploitation of this new computational environment will be enabled by fundamental computer science research on new algorithms, data structures, system software, information mining and visualization techniques, and collaborative environments for data exploration and analysis.

1. "The word 'grid' is chosen by analogy with the electric power grid, which provides pervasive access to power and, like the computer and a small number of other advances, has had a dramatic impact on human capabilities and society. We believe that by providing pervasive, dependable, consistent and inexpensive access to advanced computational capabilities, databases, sensors, and people, computational grids will have a similar transforming effect, allowing new classes of applications to emerge." From the Preface to *The Grid, Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, Inc. (1999), edited by Ian Foster and Carl Kesselman.

Cognizant Program Officer(s):

- Richard L. Hilderbrandt, Program Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-7093, e-mail: rhilderb@nsf.gov.
- Richard Hirsh, Deputy Division Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-8970, e-mail: rhirsh@nsf.gov.
- Robert Borchers, Division Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-8970, e-mail: rborcher@nsf.gov.

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

- 47.070 --- Computer and Information Science and Engineering

ELIGIBILITY INFORMATION

- **Organization Limit:** Any U.S. organization with prior experience in providing leading edge computational capability to a wide range of users is eligible to apply. Synergistic collaboration among two or more geographically distinct sites is required. Collaboration or partnerships with industry or government laboratories is encouraged where appropriate. Please see the full announcement/solicitation for further information.
- **PI Eligibility Limit:** No individual may be Principal Investigator or Co-PI on more than one proposal.
- **Limit on Number of Proposals:** None

AWARD INFORMATION

- **Anticipated Type of Award:** Cooperative Agreement
- **Estimated Number of Awards:** 1
- **Anticipated Funding Amount:** \$45 million for all major components of the distributed facility in FY01, plus up to \$4 million per year for operating costs in FY02 and FY03. Upgrade funds possible in future years.

PROPOSAL PREPARATION AND SUBMISSION INSTRUCTIONS

A. Proposal Preparation Instructions

- **Full Proposals:** Supplemental Preparation Guidelines
 - The program announcement/solicitation contains supplements to the standard Grant Proposal Guide (GPG) proposal preparation guidelines. Please see the full program announcement/solicitation for further information.

B. Budgetary Information

- **Cost Sharing Requirements:** Cost Sharing is required (Percentage).
- **Cost Sharing Level/Amount:** 20%
- **Indirect Cost (F&A) Limitations:** Not Applicable.
- **Other Budgetary Limitations:** Not Applicable.

C. Deadline/Target Dates

- **Letters of Intent (*optional*):** None
- **Preliminary Proposals (*optional*):** None
- **Full Proposal Deadline Date(s):** April 19, 2001

D. FastLane Requirements

- **FastLane Submission:** Full Proposal Required
- **FastLane Contact(s):**
 - Richard L. Hilderbrandt, Program Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-7093, e-mail: rhilderb@nsf.gov.

PROPOSAL REVIEW INFORMATION

- **Merit Review Criteria:** National Science Board approved criteria. Additional merit review considerations apply. Please see the full program announcement/solicitation for further information.

AWARD ADMINISTRATION INFORMATION

- **Award Conditions:** Additional award conditions apply. Please see the program announcement/solicitation for further information.
- **Reporting Requirements:** Standard NSF reporting requirements apply.

TABLE OF CONTENTS

SUMMARY OF PROGRAM REQUIREMENTS

I. INTRODUCTION

II. PROGRAM DESCRIPTION

III. ELIGIBILITY INFORMATION

IV. AWARD INFORMATION

V. PROPOSAL PREPARATION AND SUBMISSION INSTRUCTIONS

A. Proposal Preparation Instructions

B. Budgetary Information

C. Deadline/Target Dates

D. FastLane Requirements

VI. PROPOSAL REVIEW INFORMATION

A. NSF Proposal Review Process

B. Review Protocol and Associated Customer Service Standard

VII. AWARD ADMINISTRATION INFORMATION

A. Notification of the Award

B. Award Conditions

C. Reporting Requirements

VIII. CONTACTS FOR ADDITIONAL INFORMATION

IX. OTHER PROGRAMS OF INTEREST

I. INTRODUCTION

NSF has a long history of support for high-performance computing and networking, beginning with the Supercomputer Centers program established in 1985 and the NSFnet. In 1998 the Partnerships for Advanced Computational Infrastructure (PACI) program replaced the Supercomputer Centers program. The PACI program added emphasis on the coupling of computational and computer science in order to more effectively exploit the emerging capabilities of scalable parallel systems, high performance networking and high bandwidth, large capacity mass storage systems, in addition to putting in place a formal education, outreach and training program.

Due to the increased computational capability that is now available, computational science is currently experiencing a revolution in its ability to solve research problems. The recent demonstration of computers with speeds of a teraflop or more (10¹² floating point operations per second) has directed attention to important fundamental science and engineering problems which are not amenable to solution with current systems, but would be accessible to terascale range computation. The President's Information Technology Advisory Committee (PITAC) final report, Information Technology Research: Investing in Our Future, released on Feb. 24, 1999, states:

"If the United States is to continue as the world leader in basic research, its scientists and engineers must have access to the most powerful computers. Therefore, the Committee recommends that the Federal government continue to provide these computing systems to the research community through major, shared-facility centers. To increase long-term, fundamental research across all science and engineering disciplines, the first priority should be to increase the computing capacity of the centers that can best serve the entire research community."

On August 3, 2000 the National Science Board (NSB), NSF's governing body, authorized the three-year award for a Terascale Computing System to the Pittsburgh Supercomputing Center following a national competition. This system will begin operation in February 2001, and the full system is anticipated to reach peak performance by October 2001.

While a terascale computing system will satisfy the current needs of a large number of scientists and engineers doing modeling and simulation, modern scientific and engineering research requires more than just computational capability. Investments in large scale research instrumentation being made in such diverse fields such as astronomy, biology, earthquake engineering, environmental science, geosciences, gravitational science, and high energy physics, will not yield their full returns unless corresponding investments are made in the infrastructure needed for data analysis. Terascale computing systems and large-scale scientific instruments and sensors are now routinely creating multi-terabyte data archives. All the researchers involved encounter similar problems since computed, observed, and experimental data all require data manipulation and storage, visualization, data mining and interpretation. The rapidly increasing rate at which data are being generated and the distance between its point of generation and those who need access to information contained in the data are problems that must be faced.

The concept of a computational grid, as enunciated in the book *The Grid: A Blueprint for a New Computing Infrastructure* (<http://www.mkp.com/grids/>) and at other forums, provides a means of meeting these needs. Briefly, the Grid is the sum of networking, computing, and data storage technologies needed to create a seamless, balanced, integrated computational and collaborative environment. An unprecedented opportunity exists to take advantage of the emerging Grid technologies to create a national infrastructure that also includes digital libraries from observations, web-based portals to a large suite of computational resources, and support for remote use of scientific instruments.

In the original PACI solicitation, NSF anticipated this need by stating "The emergence of scaleable parallel systems, high performance networking and high bandwidth, large-capacity mass storage systems creates the opportunity for a national infrastructure consisting of a number of geographically distributed sites strongly coupled to high-end computational resources and to each other via high-speed communication networks". The possibility now presents itself to take further advantage of current capabilities in order to more tightly couple various distributed resources like computing systems, data repositories, visualization systems, etc., by creating a system that is part of the national computational grid and enhances the capabilities available over those located at just one site, e.g., data acquisition, control of remote instruments, and computational steering.

In order to fully exploit such systems, researchers will require new algorithms, data structures, advanced system software, distributed access to very large data archives, sophisticated information mining and visualization techniques, and collaborative environments for data exploration and analysis. Important research advances will be required in every aspect of high-performance computing, communications and information processing, necessitating a long-term, sustained, and coordinated research program. Given the greater complexity of highly parallel computer architectures, high-speed networks, petabyte data archives, and visualization and teleimmersion systems, a balanced approach to the deployment of a distributed facility that can effectively integrate a terascale system is even more critical.

The concept of a computational facility as an isolated entity, an island where everything necessary for the solution and analysis of problems exists, has been slowly changing due to the PACI partnerships and the emergence of the Grid. This program will establish an advanced, balanced, multi-site "distributed facility" connected by ultra high speed networks that will significantly enhance the capabilities of U.S. researchers in all areas of computational, computer and information science and engineering. This facility will in turn open a pathway to future, even larger scale, computing and information environments.

II. PROGRAM DESCRIPTION

The purpose of this solicitation is to continue NSF's role of serving the needs of the high end computational and computer science research community in both simulation and data intensive computation. It will provide the major infrastructure needed to lead to breakthroughs and enable further advances across all fields supported by NSF. It will also take advantage of and contribute to the emergence of the Grid as another means of providing computational capability, sophisticated data handling and interaction with remote sites.

With this in mind, NSF seeks to fund an advanced "distributed facility" that will demonstrate both single site, and "Grid enabled" capabilities for both simulation and data exploration beyond what is available at current PACI sites. The goal is to deliver production quality service from the distributed facility. The need for a distributed facility also acknowledges that the expertise for the various components of this new type of facility do not necessarily all reside in the same location, and that by collaborative efforts, great synergies can be enabled.

The sort of resources one might expect to be found in the distributed facility could be, but are not limited to, the examples given below.

- one computing system capable of five or more teraflops per second (peak) performance located at a single site,
- another large, but not necessarily comparably configured system at another site coupled with the first to test large scale distributed computing across The DTF and other resources.
- a networked system optimized to support the use of data stored at one site by a major computational resource at a geographically distant site
- visualization facilities allowing data residing at possibly more than one site to be viewed on a system remote from the data
- a distributed storage system allowing data to be stored at various sites on the Grid from a single computational resource
- ultra high-speed network connections that will enable computational resources to access unique scientific instruments directly for experiments in on-line control and data collection

The proposed distributed facility will add to the already existing capabilities provided by NSF and form the foundation of a distributed computational infrastructure that will meet the growing demands for modeling and simulation as well as anticipate the current and future needs of the scientific and engineering communities dealing with exceptionally large data intensive information management applications. The distributed facility will include substantial support for accessing, analyzing, processing, transmitting, and visualizing multi-terabyte data collections of current and future interest to the U.S. research community. This will require the DTF to have terabytes to petabytes of online and archival storage available for user access and multi-gigabit per second network connectivity. Fundamental computer science research, such as the FY00 ITR GriPhyN project, will be necessary to achieve truly effective use of this national resource.

It is the expectation of the NSF that this system will be fully coordinated with the resources and activities of the existing PACI partnerships, such as, but not limited to, user support and consulting. As such, the proposed facility will be managed for the use of the national community in cooperation with the Division of Advanced Computational Infrastructure and Research at NSF. This multi-site distributed facility will be integrated into the nationwide Grid and will supplement the capabilities that are available through the PACI partnerships.

The goal of this solicitation is to achieve the most computational infrastructure for the broadest scientific and engineering community within the funds available. Since the proposed distributed facility is not simply a terascale computing platform, the requirements of this solicitation can only be met by a collaborative proposal involving two or more geographically distinct sites. Achieving the goals set forth in this solicitation will require a joint development effort between individual sites, multiple vendors, the PACI partnerships, computer scientists, software engineers, and other possible collaborators. The problems to be resolved in implementing a system of this scale are complex. Having collaborative expertise available to address the multitude of problems that are likely to occur will enhance the chances for success.

It is anticipated that only a portion of the \$45 million need be spent on a five or more teraflop computing system located at a single site, and that the successful DTF should be a balanced system. It is also anticipated that existing equipment at the collaborating sites involved in a successful proposal may be made available to, or upgraded for integration into, the proposed distributed facility. It is therefore important to include substantial documentation on existing resources, and detailed plans on how such resources might be utilized as part of the overall project.

III. ELIGIBILITY INFORMATION

Any U.S. organization with prior experience in providing leading edge computational capability to a wide range of users is eligible to apply.

Synergistic collaboration among two or more geographically distinct sites is required. Collaboration or partnerships with industry or government laboratories is encouraged where appropriate.

No individual may be Principal Investigator or Co-PI on more than one proposal.

No proposing organization may submit or participate in more than one proposal.

IV. AWARD INFORMATION

Subject to the availability of funds, a total of \$45 million for all major components of the distributed facility, including the terascale computing system (hardware, software, communications and maintenance), in Fiscal Year 2001, plus up to \$4 million per year for operating costs in future years, is available for this competition. NSF expects to make this award in May 2001. This award will be for three years. Pending availability of funds, the system's capabilities may be upgraded in subsequent years.

V. PROPOSAL PREPARATION AND SUBMISSION INSTRUCTIONS

A. Proposal Preparation Instructions

Full Proposal:

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 Web Site at: <http://www.nsf.gov/cgi-bin/getpub?nsf012>. 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.

This solicitation requires the submission of collaborative proposals involving multiple institutions. While there is no minimum or maximum number of collaborators or partners set forth in this solicitation, we expect proposals to include only those who will be significant participants, and not simply institutions that may have expressed an interest, or will have a minimal relationship with the awardee(s). An excessive number of collaborators or partners, therefore, is strongly discouraged. See section II.C.11.b of the GPG for special instructions for submission of collaborative proposals.

Substantial documentation must be included on existing resources, and detailed plans on how such resources might be utilized as part of the overall project.

The proposal must address the three additional review criteria in section VI.A of this solicitation, in addition to the standard NSF criteria.

GPG exceptions:

- The project description section of the proposal has a limit of 25 pages.
- The budget justification section has no page limit.
- A budget for total funds requested from NSF for equipment acquisition and operations should be submitted via FastLane. Separate hard copy budgets reflecting: 1) costs for initial system acquisition in FY 2001; and 2) operations budgets for FY 2002 and 2003, should be submitted with the signed cover sheets (see Section D. FastLane Requirements).
- No limit has been set on the number of appendices (which are placed in the Supplementary Documentation section of FastLane), but only essential information should be included. Excessive appendices, which make the body of the proposal its smallest component, are strongly discouraged.

Proposers are reminded to identify the program solicitation number (NSF 01-51) in the program announcement/solicitation block on the proposal Cover Sheet (NSF Form 1207). 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 at a level of 20 percent of the requested total amount of NSF funds is required for all proposals submitted in response to this solicitation. The proposed cost sharing must be shown on Line M on the proposal budget. Documentation of the availability of cost sharing must be included in the proposal.

Only items which would be allowable under the applicable cost principles, if charged to the project, may be included in the awardee's contribution to cost sharing. Contributions may be made from any non-Federal source, including non-Federal grants or contracts, and may be cash or in kind (see OMB Circular A-110, Section 23). It should be noted that contributions counted as cost sharing toward projects of another Federal agency may not be counted towards meeting the specific cost sharing requirements of the NSF award.

All cost sharing amounts are subject to audit. Failure to provide the level of cost sharing reflected in the approved award budget may result in termination of the NSF award, disallowance of award costs and/or refund of award funds to NSF.

C. Deadline/Target Dates

Proposals must be submitted by the following date(s):

Full Proposals by 5:00 PM local time: April 19, 2001

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 1-800-673-6188.

Submission of Signed Cover Sheets. The signed copy of the proposal Cover Sheet (NSF Form 1207) must be postmarked (or contain a legible proof of mailing date assigned by the carrier) within five working days following proposal submission and be forwarded to the following address:

National Science Foundation
DIS – FastLane Cover Sheet
4201 Wilson Blvd.
Arlington, VA 22230

VI. PROPOSAL REVIEW INFORMATION

A. NSF Proposal Review Process

Reviews of proposals submitted to NSF are solicited from peers with expertise in the substantive area of the proposed research or education project. These reviewers are selected by Program Officers charged with the oversight of the review process. NSF invites the proposer to suggest at the time of submission, the names of appropriate or inappropriate reviewers. Care is taken to ensure that reviewers have no conflicts with the proposer. Special efforts are made to recruit reviewers from non-academic institutions, minority-serving institutions, or adjacent disciplines to that principally addressed in the proposal.

Proposals will be reviewed against the following general review criteria established by the National Science Board. Following each criterion are potential considerations that the reviewer may employ in the evaluation. These are suggestions and not all will apply to any given proposal. Each reviewer will be asked to address only those that are relevant to the proposal and for which he/she 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?

Principal Investigators should address the following elements in their proposal to provide reviewers with the information necessary to respond fully to both of the above-described NSF merit review criteria. NSF staff will give these elements careful consideration 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 additional, more specific, criteria dealing with the ability of the proposer to establish a distributed facility.

1. What does the DTF provide that will enable breakthrough science in all areas supported by NSF? What advanced capabilities does the overall proposed distributed facility provide to support modern computational, computer and information science and engineering? Will these capabilities be truly usable by scientists in these fields, as well as by computational specialists? Will the capabilities be integrated, so that the sum is more than a collection of the parts? This criterion has several subparts, each of which will be reviewed in all proposals received.

1a. How well does the proposed infrastructure support high-end scientific computation?

What is the expected sustained performance of the system to be placed at the single site on real problems? Do the secondary sites provide significant additional computational capability? Are all the sites involved well balanced, both within each system and between sites? Is appropriate system and application software available to support scientific computations? Will the proposed system be able to support a variety of applications representing the broad community of potential NSF users? There are no benchmark requirements mandated in this solicitation; however, any current performance data from either benchmarks or actual application codes are strongly encouraged in the proposal.

A single-site, five or more teraflop (peak) system must be a part of any successful proposal. However, the performance of this system will not be the only, nor necessarily the deciding, factor in the eventual decision. The overall capability of the DTF will be paramount.

1b. How well does the proposed infrastructure support data-intensive computation?

Will the proposed facility be able to deal effectively with extremely large datasets? How much data are accessible on a regular basis? Is the I/O and network bandwidth sufficient to support the use of these data? Will it be possible to provide parallel I/O to archival storage? Is there appropriate support for visualization and other analysis of massive data? Is appropriate software available to support the access, manipulation, visualization, and transmittal of data? Will differing storage systems/architectures at participating sites be an obstacle to data transfer/access? How will data integrity be maintained?

1c. How well does the proposed infrastructure support remote access to distributed unique resources?

Do the connections to major computation and data resources have sufficient capacity? Will the proposed facility provide the necessary networking bandwidth between sites to facilitate Grid-enabled computing and the sharing of multi-terabyte data objects? For a facility such as this networking bandwidth and latency become intertwined issues. Can the networking be configured for appropriate experiments? Are there other unique capabilities, e.g., remote operation of scientific instruments, accessible through the network that might be used for cutting edge science? Does the network middleware exist to support the required networking for scientific users?

2. What is the procedure for ensuring the DTF (and possible upgrades) will be available for use by the national research community?

Is the plan for implementation and support of the infrastructure (both the five or more teraflop computational engine and the other systems proposed) viable? Is the proposed infrastructure sufficiently scalable and what is (are) the future upgrade path(s) available to continue to improve performance for the next three to five years? What are the qualifications/experience of the PI and the proposing organization in regard to managing a set of resources for national use? How does the proposed infrastructure relate to those at other sites in delivering high-end services to the national research community?

How will the new facility be integrated into the existing PACI program to both take advantage of existing expertise and to minimize operating costs? Does the proposal describe plans for current and future physical facilities that will house the components of the distributed facility including schedule implications of providing computer-ready space? Does the proposal adequately address delivery and upgrade dates and acceptance milestones?

3. What is the cost effectiveness to NSF of the proposed system, including operations?

What is the cost effectiveness of the planned system and its operations? How will the required cost sharing be applied? Overall cost effectiveness of the proposed activity is an important consideration for all proposals submitted under this solicitation.

A summary rating and accompanying narrative will be completed and signed 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 Director. In addition, the proposer will receive an explanation of the decision to award or decline funding.

B. Review Protocol and Associated Customer Service Standard

All proposals are carefully reviewed by at least three other persons outside NSF who are experts in the particular field represented by the proposal. External panels using the criteria identified in this solicitation will review all proposals. Panelists will be chosen from academia, the private sector, the national laboratories, and other government agencies. They will have expertise in all areas of high-performance computing necessary to evaluate the proposals received, including computational science, computer science, networking, data, and scientific disciplines representative of the target user community. Organizations submitting proposals deemed to have sufficient merit may be site visited and, if this is done, the resulting site visit reports will be a critical part of the deliberations of the final panel.

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.

In most cases, proposers will be contacted by the Program Officer after his or her recommendation to award or decline funding has been approved by the Division Director. This informal notification is not a guarantee of an eventual award.

NSF will be able to tell applicants whether their proposals have been declined or recommended for funding within six months for 95 percent of proposals. The time interval begins on the proposal deadline or target date or from the date of receipt, if deadlines or target dates are not used by the program. The interval ends when the Division Director accepts the Program Officer's recommendation.

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 its 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 Web site at http://www.nsf.gov/home/grants/grants_gac.htm. Paper copies may be obtained from the NSF Publications Clearinghouse, telephone (793) 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 Web site 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 Web site at <http://www.gpo.gov>.

Special Award Conditions

The terascale environment will be managed by the Awardees in cooperation with NSF in the context of the Advanced Computational Infrastructure (ACI) program. The five teraflop system to be put in place represents a high end capability provided by the ACI Program, but will be coordinated with all other resources available through ACI. Other parts of the proposed infrastructure will also be coordinated with capabilities provided by the PACI partnerships. For example, the existing National Resource Allocations Committee (NRAC) that is responsible for large allocations on current PACI systems will manage the allocation of the usage on the five or more teraflop system. NRAC may also manage the allocation of other computational and data resources acquired as part of the proposed infrastructure. Title to the equipment purchased as a result of this solicitation vests in the Awardees. The NSF may, at any time prior to 120 calendar days after the expiration date of the Cooperative Agreement, direct the Awardee to transfer title to the equipment to the Federal Government or a third party named by the Federal Government.

C. Reporting Requirements

For all multi-year grants (including both standard and continuing grants), the PI must submit an annual project report to the cognizant Program Officer at least 90 days before the end of the current budget period.

Within 90 days after the expiration of an award, the PI also is required to submit a final project report. Approximately 30 days before expiration, NSF will send a notice to remind the PI of the requirement to file the final project report. Failure to provide final technical reports delays NSF review and processing of pending proposals for that PI. PIs should examine the formats of the required reports in advance to assure availability of required data.

NSF has implemented an electronic project reporting system, available through FastLane. This system permits electronic submission and updating of project reports, including information on project participants (individual and organizational), activities and findings, 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.

VIII. CONTACTS FOR ADDITIONAL INFORMATION

General inquiries regarding Distributed Terascale Facility should be made to:

- Richard L. Hilderbrandt, Program Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-7093, e-mail: rhilderb@nsf.gov.
- Richard Hirsh, Deputy Division Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-8970, e-mail: rhirsh@nsf.gov.
- Robert Borchers, Division Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-8970, e-mail: rborcher@nsf.gov.

For questions related to the use of FastLane, contact:

- Richard L. Hilderbrandt, Program Director, CISE/ACIR, Advanced Computational Infrastructure and Research, Rm. 1122, telephone: (703)292-7093, e-mail: rhilderb@nsf.gov.

IX. OTHER PROGRAMS OF INTEREST

The NSF *Guide to Programs* is a compilation of funding for research and education in science, mathematics, and engineering. The NSF *Guide to Programs* is available electronically at <http://www.nsf.gov/cgi-bin/getpub?gp>. General descriptions of NSF programs, research areas, and eligibility information for proposal submission are provided in each chapter.

Many NSF programs offer announcements or solicitations concerning specific proposal requirements. To obtain additional information about these requirements, contact the appropriate NSF program offices. Any changes in NSF's fiscal year programs occurring after press time for the *Guide to Programs* will be announced in the NSF E-Bulletin, which is updated daily on the NSF web site at <http://www.nsf.gov/home/ebulletin>, and in individual program announcements/solicitations. Subscribers can also sign up for NSF's Custom News Service (<http://www.nsf.gov/home/cns/start.htm>) to be notified of new funding opportunities that become available.

ABOUT THE NATIONAL SCIENCE FOUNDATION

The National Science Foundation (NSF) funds research and education in most fields of science and engineering. Awardees are wholly responsible for conducting their project activities and preparing the results for publication. Thus, the Foundation does not assume responsibility for such findings or their interpretation.

NSF welcomes proposals from all qualified scientists, engineers and educators. The Foundation strongly encourages women, minorities and persons with disabilities to compete fully in its programs. In accordance with Federal statutes, regulations and NSF policies, no person on grounds of race, color, age, sex, national origin or disability shall be excluded from participation in, be denied the benefits of, or be subjected to discrimination under any program or activity receiving financial assistance from NSF (unless otherwise specified in the eligibility requirements for a particular program).

Facilitation Awards for Scientists and Engineers with Disabilities (FASSED) provide funding for special assistance or equipment to enable persons with disabilities (investigators and other staff, including student research assistants) to work on NSF-supported projects. See the program announcement/solicitation for further information.

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 or (800) 281-8749, FIRS at 1-800-877-8339.

The National Science Foundation is committed to making all of the information we publish easy to understand. If you have a suggestion about how to improve the clarity of this document or other NSF-published materials, please contact us at plainlanguage@nsf.gov.

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; 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 applicant 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 needing information as part of the review process or in order to coordinate programs; 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," 63 Federal Register 267 (January 5, 1998), and NSF-51, "Reviewer/Proposal File and Associated Records," 63 Federal Register 268 (January 5, 1998). Submission of the information is voluntary. Failure to provide full and complete information, however, may reduce the possibility of receiving an award.

Pursuant to 5 CFR 1320.5(b), an agency may not conduct or sponsor, and a person is not required to respond to an information collection unless it displays a valid 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 this burden estimate and any other aspect of this collection of information, including suggestions for reducing this burden, to: Suzanne Plimpton, Reports Clearance Officer, Information Dissemination Branch, Division of Administrative Services, National Science Foundation, Arlington, VA 22230, or to Office of Information and Regulatory Affairs of OMB, Attention: Desk Officer for National Science Foundation (3145-0058), 725 17th Street, N.W. Room 10235, Washington, D.C. 20503.

NSF 01-51

OMB control number: 3145-0058.