

**Office of Science  
Financial Assistance  
Funding Opportunity Announcement  
DE-PS02-08ER08-19**

***Software Development Tools for  
Improved Ease-of-Use of Petascale Systems***

The Office of Advanced Scientific Computing Research (ASCR) of the Office of Science (SC), U.S. Department of Energy (DOE), hereby announces its interest in receiving applications for research grants in software development tools for improved ease-of-use of petascale systems.

Petascale computing systems soon will be available to the DOE science community. Such systems will exhibit increased architectural complexity and tens to hundreds of thousands of processor cores. Increased architectural complexity includes multicore/heterogeneous CPUs, novel memory systems and intelligent interconnects. Applications are also becoming more complex with a variety of languages, libraries, programming models, data structures, and algorithms in a single application. Taken together, these trends generate a critical need for tools that can help application teams address severe complexity and scalability challenges.

Software development tools serve as a key interface between application teams and target HPC architectures. Broadly speaking, tool functionality can be decomposed into three categories: correctness tools which support the rapid debugging of complex code, performance tools for identifying and removing performance bottlenecks, and development environments which enable the efficient generation and test of complex codes and code frameworks. Both correctness and performance tools must be fully scalable in order to address subtle problems that may be manifested only at large scale, and they must rely on scalable infrastructures that support tool communication, data management, binary manipulation of application executables, and a variety of other capabilities.

This announcement is focused on research and development for innovations in petascale tools in each of these areas: correctness tools, performance tools, scalable tool infrastructure and development environments. The activities supported by this notice may be a combination of basic and applied research, development, prototyping, testing and ultimately deployment. Partnerships among universities, National Laboratories, and industry are encouraged.

**LETTERS OF INTENT: May 12, 2008. 5:00 p.m., Eastern Time**

**Potential applicants are REQUIRED to submit a two-page Letter of Intent (LOI) by email to [petascaletools@ascr.doe.gov](mailto:petascaletools@ascr.doe.gov). The subject line of the email should be: "Letter of Intent for Announcement DE-PS02-08ER08-19". The LOI should be a Word file attached to the email. No FAX or mail submission of Letters of Intent will be accepted. Letters of Intent must be received by May 12, 2008, 5:00 p.m., Eastern Time.**

The purpose of a LOI is to save the time and effort of applicants in preparing and submitting a formal project application that may be inappropriate for the program. Letters of Intent also assist ASCR in planning the peer review process and the selection of properly qualified reviewers.

Letters of Intent should consist of no more than two pages total. The LOI should provide (1) the Principal Investigator's name, telephone number, and email address; (2) the name of the Principal Investigator's employing institution; (3) the title of the proposed research; (4) a clear and concise description of the proposed research and research objectives; (5) a statement of background and significance of the proposed project; (6) a rough dollar approximation of the budget for each year of the proposed research; (7) a curriculum vita that highlights the Principal Investigator's expertise and background in successful research related to the subject of this announcement and the proposed research; and (8) the proposed research team and brief statements of their expertise. A Word form for the LOI is available at:

<http://www.science.doe.gov/ascr/Research/08CSSolicit.html>, submitters are strongly encouraged to use this form for their LOI submission.

Letters of Intent will be reviewed for conformance with the guidelines and technical areas provided in this announcement. **A response to a LOI encouraging or discouraging formal applications will be communicated to all applicants by May 26, 2008.** Applicants who have not received a response regarding the status of their LOI by this date are responsible for contacting the program to confirm their status. **Formal applications will be accepted only from those encouraged to submit a formal application in response to their LOI. No other formal applications will be considered.**

**APPLICATION DUE DATE: July 18, 2008, 8:00 p.m., Eastern Time**

Applications must be submitted using [Grants.gov](http://Grants.gov), the Funding Opportunity Announcement can be found using the CFDA Number, 81.049 or the Funding Opportunity Announcement number, DE-PS02-08ER08-19. Applicants must follow the instructions and use the forms provided on [Grants.gov](http://Grants.gov).

**FOR FURTHER INFORMATION CONTACT:**

Dr. Frederick Johnson  
Telephone: (301) 903-5800  
Fax: (301) 903-7774  
E-mail: [fjohnson@ascr.doe.gov](mailto:fjohnson@ascr.doe.gov)

**SUPPLEMENTARY INFORMATION:**

Software development tools enable application teams to effectively use large scale systems for the efficient execution of complex scientific applications. They are essential to the success of both large scale systems and complex applications. Next generation petascale systems will have tens to hundreds of thousands of processors, an unprecedented level of complexity, and will require significant new levels of scalability and functionality in software tools. A new and innovative generation of software development tools is needed to meet and surpass application requirements for scalability, functionality, reliability, and ease of use.

The complexity and scale of petascale systems and large application codes represent major challenges for development tools including: radical increases in node and processor core counts, support for multi-mode parallelism, reduced memory per core, heterogeneous nodes, and support for fault tolerance. Application developer and user needs for these systems include: a means for debugging at scale, increased support for memory debugging, memory characterization tools, both lightweight and heavyweight tools, performance analysis tool support for serial code segments, multithreaded segments and multimode segments, and means for understanding and optimizing for topology related performance.

The research activities supported by this activity need to bridge the gap between large complex applications and next-generation hardware, including interactions with novel architectures. Consequently, there are a wide variety of research topics that are appropriate for this effort. Example candidate topics are provided below, but research in other relevant areas and combinations of areas is encouraged.

### **Performance Tools**

**Automated Diagnosis and Remediation** -- New approaches to performance optimization which move beyond manual methods and enable greater automation and which support automation of diagnosis, optimization and anomaly detection.

**Load Imbalance Detection** - Highly scalable methodology for detecting load imbalance in applications running on hundreds of thousands of processor cores. Tools which provide root cause analysis in addition to detection.

**Heterogeneous, Hierarchical Architecture Support** -- Performance tools which support multilevel parallel paradigms, including hybrid OpenMP/MPI programs. Tools which capture and relate performance and reliability problems to source code in ways that make multilevel performance optimization possible and practical.

### **Correctness Tools**

**Scalable Debuggers** - Both lightweight and heavyweight approaches to scalable debugging that support ease of use, error detection at scale, and in-depth root cause analyses.

**Memory Usage** - Both lightweight and heavyweight tools for monitoring memory utilization (especially memory leaks and overall memory consumption) and tools to find programming errors in the way memory is accessed.

**Thread Correctness** -- Multi-platform tools that enable users to detect incorrect use of parallel programming techniques including thread correctness checkers and Message Passing Interface (MPI) usage checkers. Tools which assess the validity of memory references, track locks that are held when memory is accessed and verify that no potential race condition exists.

### **Scalable Infrastructure**

**Data Management and Communication** -- Support for all aspects of the gathering, reduction, and storage of application information and metadata. Support for communicating information among tool components on different nodes, getting information from external sources such as the operating system, compiler, scheduler, and runtime system, and exchanging information between tools.

**Scheduler and Operating System Interaction** - Support for close coordination of tools with the scheduler, e.g. for tool launch on multiple nodes, and the operating system, e.g. process control interfaces for access to thread information and low overhead access to hardware counters.

**Binary Manipulation** - Support for binary analysis of optimized and stripped programs, and the ability to generate new binaries with instrumentation.

## **Development Environment**

**Application Build Tools** - Tool support for radical improvements in the management of the application build process that address the complexities arising from multiple target systems, operating systems, libraries and software versions. Also support for common option sets, command line interfaces, shared libraries and dynamic link order.

**Mixed Language Environments** - Tool support for mixed language programming including traditional languages, Fortran, C, C++; scripting languages, Python; and emerging languages such as the PGAS languages UPC, Co-array Fortran and the HPCS languages.

**Compiler Infrastructure** -- A flexible, extendible, portable, open source compiler infrastructure to support efficient information transfer between compile time analyses and tools and runtime analyses and tools.

**Program Transformations** - Tools supporting source-to-source transformations to enable codes to automatically adapt to new computer architectures achieve maximum architecture independence and efficiently use complex libraries.

**Integrated Development Environments (IDEs)** - Integrated frameworks supporting the effective integration of development and runtime environments and achieve significant improvements in programmer productivity in the creation of complex application codes.

## **References**

These example research topics represent only a portion of the research challenges for petascale tools. All interested proposers are strongly encouraged to study the following references for additional discussion insight:

Software Development Tools for Petascale Computing Workshop Presentations:  
<http://www.csm.ornl.gov/workshops/Petascale07/presentations.html>

Software Development Tools for Petascale Computing Workshop Final Report:  
[http://www.csm.ornl.gov/workshops/Petascale07/sdtpc\\_workshop\\_report.pdf](http://www.csm.ornl.gov/workshops/Petascale07/sdtpc_workshop_report.pdf)

## **Community Building**

An important goal of this notice is to foster active, integrated research community in petascale tools for high end systems. Consequently the following are mandatory requirements for awardees:

- All developed code must be released under the most permissive open source license possible. This is to enable other researchers and vendors to build upon research successes with a minimum of intellectual property issues.
- Each research team should plan to send representatives to annual PI meetings and give presentations on the status and promise of their research. Meeting attendees will include invited participants from other relevant research communities. The objectives of these meetings include fostering a sense of community and serving as a venue for exchange of information with complementary programs including the DARPA HPCS program, NSF programs in CISE and OCI, NNSA ASC program, and the DOE/SC SciDAC program.

## **Testbed Access**

Applications should provide a plan for utilizing leadership class systems at Oak Ridge National Laboratory and Argonne National Laboratory and systems at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory for the purpose of software testing at scale. Each application should contain a section which discusses the characteristics of the test environments necessary for the research and identify the time frames in which specific testbed support will be required. Since relatively limited amounts of testing time will be available on these systems, the individual testing plans will be used to develop an overall test plan for the program.

## **Program Funding**

It is anticipated that up to \$3 million annually will be available for multiple awards for this program. Awards are planned to be made in Fiscal Year 2009, and applications may request project support for up to three years. Annual budgets for successful projects are expected to range from \$250,000 to \$700,000 per project although smaller projects of exceptional merit may be considered. Annual budgets may increase in the out-years but should remain within the overall annual maximum guidance. All awards are contingent on the availability of funds and programmatic needs. DOE is under no obligation to pay for any costs associated with the preparation or submission of an application. DOE reserves the right to fund, in whole or part, any, all, or none of the applications submitted in response to this Notice.

## **Merit Review Criteria**

Applications will be subjected to scientific merit review (peer review) and will be evaluated against the following evaluation criteria which are listed in descending order of importance codified at 10 CFR 605.10(d):

1. Scientific and/or Technical Merit of the Project;
2. Appropriateness of the Proposed Method or Approach;
3. Competency of Applicant's Personnel and Adequacy of Proposed Resources; and
4. Reasonableness and Appropriateness of the Proposed Budget.

The evaluation process will include program policy factors, such as the relevance of the proposed research to the terms of the solicitation and the agencies' programmatic needs. Note that external peer reviewers are selected with regard to both their scientific expertise and the absence of conflict-of-interest issues. Both Federal and non-Federal reviewers may be used, and submission of an application constitutes agreement that this is acceptable to the investigator(s) and the submitting institution.

Posted on the Office of Science Grants and Contracts Web Site  
April 16, 2008.