

**Program Announcement  
To DOE National Laboratories  
LAB 07-23**

***Operating and Runtime Systems  
for Extreme Scale Scientific Computation***

**SUMMARY:** 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 proposals for research in Operating and Runtime Systems for Extreme Scale Scientific Computation (FASTOS). This announcement is focused on research and development of operating and runtime systems which enable the effective management and use of extreme-scale systems (petascale and beyond) for scientific computation. The overall goal of this announcement is to stimulate research and development related to operating and runtime systems for petascale systems the in 2010 to 2015 timeframe. It is likely that these systems will include a combination of commodity and custom components, with different systems reflecting different degrees of customization. Operating and runtime systems research must be driven from the needs of current and future applications, and the primary focus is on supporting the needs of existing and anticipated SC and other DOE applications. An ultimate goal would be the development of a unified operating and runtime system that could fully support and exploit petascale and beyond systems and autonomously adapt to meet specific application needs for performance, functionality, security, and fault tolerance. The activities supported by this notice may be a combination of basic research, development, prototyping, and testing. Partnerships among universities, National Laboratories, and industry are encouraged.

**PREPROPOSAL DUE DATE:** April 6, 2007, 4:30 pm, Eastern Time

Potential researchers are **required** to submit a two-page preproposal by email to [fjohnsonr@ascr.doe.gov](mailto:fjohnsonr@ascr.doe.gov). Preproposals must be received by **April 6, 2007, 4:30 p.m., Eastern Time**. The subject line of the email should be: "FASTOS Preproposal". The preproposal should be a Word file attached to the email, having 1 inch margins when printed. No FAX or mail submission of preproposals will be accepted.

Preproposals will be reviewed for conformance with the guidelines and technical areas specified in this announcement. A response to preproposals encouraging or discouraging formal proposals will be communicated to the proposers by April 13, 2007. Proposers who have not received a response regarding the status of their preproposal by this date are responsible for contacting the program to confirm their status.

Preproposals should consist of no more than two pages total. This narrative should give the project title and describe the research objectives, the technical approach(s), and all proposed team members and their expertise. It should also include a rough estimate of the planned budget request. The intent in requesting a preproposal is to save the time and effort of researchers in preparing and submitting a formal project proposal that may be inappropriate for the program.

Preproposals also assist ASCR in planning the peer review process and the selection of potential reviewers for the proposal. **Formal proposals will be accepted only from preproposers encouraged to submit a formal proposal.**

**PROPOSAL DUE DATE:** June 11, 2007, 8:00 pm, Eastern Time

Full proposals submitted in response to this Announcement must be submitted to the DOE Electronic Proposal Management Application (ePMA) system (<https://epma.doe.gov>) no later than 8:00 p.m., Eastern Time, June 11, 2007, to be accepted for merit review and to permit timely consideration for award in Fiscal Year 2008. It is important that the entire peer reviewable proposal be submitted to the ePMA system as a single PDF file attachment.

In order to expedite the review process, it is **essential to also submit via email a single PDF file of the entire LAB proposal and FWP** addressed to Dr. Frederick Johnson at: [fjohnson@ascr.doe.gov](mailto:fjohnson@ascr.doe.gov). Please use "FASTOS Proposal" as the subject of the email.

To identify that the FWP is responding to this program announcement, please fill in the following fields in the "ePMA Create Proposal Admin Information" screen as shown:

**Proposal Short Name:**

**Fiscal Year:**

**Proposal Reason:**

**Program Announcement Number:** LAB 07-23\*

**Program announcement Title:** Operating and Runtime Systems for Extreme Scale Scientific Computation\*

**Proposal Purpose:**

**Estimated Proposal Begin Date:**

**HQ Program Manager Organization:**

**\* Please use the wording shown when filling in these fields to identify that the FWP is responding to this Program Announcement.**

DOE National Laboratories should submit using ePMA as instructed above. Researchers from other Federal agencies and Non-DOE Federally Funded Research and Development Centers (FFRDCs) should follow the format at [http://www.science.doe.gov/grants/fed\\_prop.html](http://www.science.doe.gov/grants/fed_prop.html) and submit via email as stated above.

**FOR FURTHER INFORMATION CONTACT:**

Dr. Frederick Johnson

**Telephone:** (301) 903-5800

**Fax:** (301) 903-7774

**E-mail:** [fjohnson@er.doe.gov](mailto:fjohnson@er.doe.gov).

**SUPPLEMENTARY INFORMATION:**

Operating and runtime systems provide mechanisms to manage system hardware and software resources for the efficient execution of large scale scientific applications. They are essential to the success of both large scale systems and complex applications. By the end of this decade petascale computers with thousands of times more computational power than any in current use will be vital tools for expanding the frontiers of science and for addressing vital National priorities. These systems will have tens to hundreds of thousands of processors, an unprecedented level of complexity, and will require significant new levels of scalability and fault management. The overwhelming size and complexity of such systems poses deep technical challenges that must be overcome to fully exploit their potential for scientific discovery. Applications require multiple services from OS/R layers, including: resource management and scheduling, fault-management (detection, prediction, recovery, and reconfiguration), configuration management, and file systems access and management. Current and future large scale parallel systems require that such services be implemented in a fast and scalable manner so that the OS/R does not become a performance bottleneck. The current trend in large scale scientific systems is to leverage operating systems developed for other areas of computing - operating systems that were not specifically designed for large scale, parallel computing platforms. Unix, Linux and other Unix derivatives are the most popular OS's in use for high end scientific computing, and these all reflect a technological heritage nearly 30-years old with few fundamental mechanisms to support parallel systems.

### **Example Research Topics**

Operating and runtime systems provide the glue that bind running applications to hardware. The research activities supported by this activity need to bridge the gap between new languages and/or programming models and next-generation hardware, including interactions with novel architectures. Consequently, there are a wide variety of research topics that are appropriate for this effort. A brief listing of candidate topics is provided below, but research in other relevant areas and combinations of areas is encouraged:

**Virtualization.** Virtualization is expected to play an increasingly important role in the deployment of large scale systems, enabling multiple operating systems on a single platform and application specific operating systems. Virtualization includes the development and use of hypervisors, virtual machine monitors, and application/runtime virtualization for HPC systems. Specific topics of interest include: identification and quantification of problems with current hypervisors in HPC systems, novel uses of hypervisors in HPC systems (development, porting, etc), support for fault handling, better support for custom hardware, and lightweight mechanisms for virtual resources.

**Fault Handling.** As the number of components in a system increase from tens to hundreds of thousands, these systems will have significantly reduced mean time between interrupt (MTI). Mechanisms to support application resiliency in the face of hardware faults are needed to support long running applications. Specific topics of interest include: tradeoffs associated with handling failures at different layers (application, runtime, OS); understanding and identifying sources of faults; approaches to proactive fault handling; fault tolerance for alternate (non- MPI) programming models; languages/APIs for the bi-directional communication of fault information between layers (e.g., between the application and runtime layers); quantification of scalability

issues; automatic, transparent, and efficient checkpoint/restart; and checkpointing when disks are far away.

**OS Noise/Interference.** Operating system interference or noise due to asynchronous overhead needed to implement system services, has been shown to have a significant impact on application performance on very large scale systems. Measurement and understanding the impact of OS interference on application performance at scale will be critical to the successful deployment of very large scale systems. Specific topics of interest include: OS design strategies for dealing with OS noise (e.g., implementations of critical services that minimize related noise and alternatives for timeouts and/or periodic service requirements); hardware features to control the impact of noise (e.g., hardware support for low overhead barriers); strategies to mitigate the impact of OS noise (e.g., exploiting asynchrony).

**Exposing Resources.** Bidirectional APIs to expose system information (performance counters) and to select implementations are critical for application level adaptability (need information about what is being used and may need to select alternate implementations). Specific topics of interest include: hooks for controlling resources; interfaces to allow code to query hardware characteristics; exposing communication related resources.

**Resource Management.** Managing the local and global resources provided by a computing system is a fundamental responsibility of any operating system, and exploration of policies and mechanisms for resource management is especially critical for petascale systems. Specific topics of interest include: local resource management (memory management, processor scheduling (multi-core), and communication support); interfaces between local and external components (gang scheduling, virtual memory reservations and queries); support for alternate (non-MPI) programming models (e.g., UPC); OS service coordination (load balancing at scale, global memory management, topology aware mapping of work- and data-units); heterogeneous resource management (HW and SW); and power management.

**Adaptability.** The ability of runtime and operating systems to change their behaviors based on application needs to improve performance or tolerate faults needed to support the use of petascale systems. Specific topics of interest include: measurement and strategies to support adaptation; understanding and exploiting application phases; adapting collective communication components; and APIs to expose resource performance models and information

**Performance Measurement.** Petascale systems will require models and tools to measure system performance, including hooks for application level performance monitoring; tools to measure runtime/OS performance; performance models (define what needs to be measured); and scalability

**System Management/Administration.** Several issues related to overall system administration need to be addressed, including: usage models (space/time sharing); flexible space-sharing; changing processors allocated to running jobs; single system image issues to ease system management number of system administrators should not scale with the size of the system; node allocation; power management; software distribution; and RAS and RAS interfaces

**Parallel I/O:** Efficient communication with external storage servers and parallel file systems is an essential component of a petascale system. Topics of interest include: support for high performance access to external servers, efficient, scalable I/O call forwarding, portable I/O models which support diverse storage instantiations, and parallel file systems.

## **Community building**

An important goal of this notice is to foster the development of an active research community in operating systems and runtime environments for high end systems. In order to meet this goal 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 or semi-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, including the Linux community. Objectives of these meetings are to foster a sense of community and serve as a venue for exchange of information. These meetings will also serve as a means to exchange information on complementary programs including the DARPA HPCS program, NNSA ASC program and DOE/SC SciDAC program.

## **Testbed access**

Proposals should provide a plan for utilizing leadership class systems at Oak Ridge National Laboratory and Argonne National Laboratory and to systems at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory for the purpose of software testing at scale. Each proposal 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. Only a relatively limited amount of testing time will be available on these systems, and the individual testing plans will be used to develop an overall test plan for the FASTOS program.

## **Program Funding**

It is anticipated that up to \$4 million annually will be available for multiple awards for this program. Awards are planned to be made in Fiscal Year 2008, and proposals may request project support for up to three years. All awards are contingent on the availability of funds and programmatic needs. Annual budgets for successful projects are expected to range from \$500,000 to \$1,000,000 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.

## **References**

FASTOS forum: <http://www.cs.unm.edu/~fastos>

Federal Plan for High-End Computing:

[http://www.nitrd.gov/pubs/2004\\_hecrtf/20040702\\_hecrtf.pdf](http://www.nitrd.gov/pubs/2004_hecrtf/20040702_hecrtf.pdf)

**OFFICE OF SCIENCE**  
**GUIDE FOR PREPARATION OF SCIENTIFIC/TECHNICAL PROPOSALS**  
**TO BE SUBMITTED BY NATIONAL LABORATORIES**

Proposals from National Laboratories submitted to the Office of Science (SC) as a result of this program announcement will follow the Department of Energy Field Work Proposal process with additional information requested to allow for scientific/technical merit review. The following guidelines for content and format are intended to facilitate an understanding of the requirements necessary for SC to conduct a merit review of a proposal. Please follow the guidelines carefully, as deviations could be cause for declination of a proposal without merit review.

### **1. Evaluation Criteria**

Proposals will be subjected to formal merit review (peer review) and will be evaluated against the following criteria which are listed in descending order of importance:

1. Scientific and/or technical merit of the project;
2. Appropriateness of the proposed method or approach;
3. Competency of proposer's personnel and adequacy of proposed resources; and
4. Reasonableness and appropriateness of the proposed budget.

The evaluation will include program policy factors such as the relevance of the proposed research to the terms of the announcement, the Department's programmatic needs, and quality of previous performance. External peer reviewers are selected with regard to both their scientific expertise and the absence of conflict-of-interest issues. Non-federal reviewers may be used, and submission of a proposal constitutes agreement that this is acceptable to the investigator(s) and the submitting institution. Proposals found to be scientifically meritorious and programmatically relevant will be selected in consultation with DOE selecting officials depending upon availability of funds in the DOE budget. The selected projects will be required to acknowledge support by DOE in all public communications of the research results.

### **2. Summary of Proposal Contents**

- Field Work Proposal (FWP) Format (Reference DOE O 412.1A) (DOE ONLY)
- Proposal Cover Page
- Table of Contents
- Budget (DOE Form 4620.1) and Budget Explanation
- Abstract (one page)
- Narrative (main technical portion of the proposal, including background/introduction, proposed research and methods, timetable of activities, and responsibilities of key project personnel)
- Literature Cited
- Biographical Sketch(es)

- Description of Facilities and Resources
- Other Support of Investigator(s)
- Appendix (optional)

## **2.1 Number of Copies to Submit**

A complete formal FWP in a single Portable Document Format (PDF) file must be submitted through the DOE ePMA system ( <https://epma.doe.gov>) as an attachment. To identify that the FWP is responding to this program announcement, please fill in the following fields in the "ePMA Create Proposal Admin Information" screen as shown:

**Proposal Short Name:**

**Fiscal Year:**

**Proposal Reason:**

**Program Announcement Number:** LAB 07-23\*

**Program announcement Title:** Operating and Runtime Systems for Extreme Scale Scientific Computation\*

**Proposal Purpose:**

**Estimated Proposal Begin Date:**

**HQ Program Manager Organization:**

**\* Please use the wording shown when filling in these fields to identify that the FWP is responding to this Program Announcement.**

In order to expedite the review process, please submit via email a single PDF file of the entire LAB proposal and FWP. The email should be addressed to Dr. Frederick Johnson at: [fjohnson@ascr.doe.gov](mailto:fjohnson@ascr.doe.gov). Please use "FASTOS Proposal" as the subject of the email.

## **3. Detailed Contents of the Proposal**

Adherence to type size and line spacing requirements is necessary for several reasons. No researcher should have the advantage, or by using small type, of providing more text in their proposals. Small type may also make it difficult for reviewers to read the proposal. Proposals must have 1-inch margins at the top, bottom, and on each side. Type sizes must be at least 11 point. Line spacing is at the discretion of the researcher but there must be no more than 6 lines per vertical inch of text. Pages should be standard 8 1/2" x 11" (or metric A4, i.e., 210 mm x 297 mm).

### **3.1 Field Work Proposal Format (Reference DOE O 412.1A) (DOE ONLY)**

The Field Work Proposal (FWP) is to be prepared and submitted consistent with policies of the investigator's laboratory and the local DOE Operations Office. Additional information is also requested to allow for scientific/technical merit review.

Laboratories may submit proposals directly to the SC Program office listed above. A copy should also be provided to the appropriate DOE operations office.

### **3.2 Proposal Cover Page**

The following proposal cover page information may be placed on plain paper. No form is required.

Title of proposed project  
SC Program announcement title  
Name of laboratory  
Name of principal investigator (PI)  
Position title of PI  
Mailing address of PI  
Telephone of PI  
Fax number of PI  
Electronic mail address of PI  
Name of official signing for laboratory\*  
Title of official  
Fax number of official  
Telephone of official  
Electronic mail address of official  
Requested funding for each year; total request  
Use of human subjects in proposed project:  
    If activities involving human subjects are not planned at any time during the proposed project period, state "No"; otherwise state "Yes", provide the IRB Approval date and Assurance of Compliance Number and include all necessary information with the proposal should human subjects be involved.  
Use of vertebrate animals in proposed project:  
    If activities involving vertebrate animals are not planned at any time during this project, state "No"; otherwise state "Yes" and provide the IACUC Approval date and Animal Welfare Assurance number from NIH and include all necessary information with the proposal.  
Signature of PI, date of signature  
Signature of official, date of signature\*

\*The signature certifies that personnel and facilities are available as stated in the proposal, if the project is funded.

### **3.3 Table of Contents**

Provide the initial page number for each of the sections of the proposal. Number pages consecutively at the bottom of each page throughout the proposal. Start each major section at the top of a new page. Do not use unnumbered pages and do not use suffices, such as 5a, 5b.

### **3.4 Budget and Budget Explanation**

A detailed budget is required for the entire project period and for each fiscal year. It is preferred that DOE's budget page, Form 4620.1 be used for providing budget information\*. Modifications

of categories are permissible to comply with institutional practices, for example with regard to overhead costs.

A written justification of each budget item is to follow the budget pages. For personnel this should take the form of a one-sentence statement of the role of the person in the project. Provide a detailed justification of the need for each item of permanent equipment. Explain each of the other direct costs in sufficient detail for reviewers to be able to judge the appropriateness of the amount requested. Further instructions regarding the budget are given in section 4 of this guide.

\* Form 4620.1 is available at web site: <http://www.science.doe.gov/grants/budgetform.pdf>

### **3.5 Abstract**

Provide an abstract of less than 400 words. Give the project objectives (in broad scientific terms), the approach to be used, and what the research is intended to accomplish. State the hypotheses to be tested (if any). At the top of the abstract give the project title, names of all the investigators and their institutions, and contact information for the principal investigator, including e-mail address.

**3.6 Narrative** (main technical portion of the proposal, including background/introduction, proposed research and methods, timetable of activities, and responsibilities of key project personnel).

The narrative comprises the research plan for the project and is limited to **20 pages (maximum)**. It should contain enough background material in the Introduction, including review of the relevant literature, to demonstrate sufficient knowledge of the state of the science. The major part of the narrative should be devoted to a description and justification of the proposed project, including details of the methods to be used. It should also include a timeline for the major activities of the proposed project, and should indicate which project personnel will be responsible for which activities.

If any portion of the project is to be done in collaboration with another institution (or institutions), provide information on the institution(s) and what part of the project it will carry out. Further information on any such arrangements is to be given in the sections "Budget and Budget Explanation", "Biographical Sketches", and "Description of Facilities and Resources".

### **3.7 Literature Cited**

Give full bibliographic entries for each publication cited in the narrative.

### **3.8 Biographical Sketches**

This information is required for senior personnel at the institution submitting the proposal and at all subcontracting institutions (if any). The biographical sketch is limited to a maximum of **two pages** for each investigator.

To assist in the identification of potential conflicts of interest or bias in the selection of reviewers, the following information **must be provided in each biographical sketch**.

**Collaborators and Co-editors:** A list of all persons in alphabetical order (including their current organizational affiliations) who are currently, or who have been, collaborators or co-authors with the investigator on a research project, book or book article, report, abstract, or paper during the 48 months preceding the submission of the proposal. Also, include those individuals who are currently or have been co-editors of a special issue of a journal, compendium, or conference proceedings during the 24 months preceding the submission of the proposal. If there are no collaborators or co-editors to report, this should be so indicated.

**Graduate and Postdoctoral Advisors and Advisees:** A list of the names of the individual's own graduate advisor(s) and principal postdoctoral sponsor(s), and their current organizational affiliations. A list of the names of the individual's graduate students and postdoctoral associates during the past five years, and their current organizational affiliations.

### **3.9 Description of Facilities and Resources**

Facilities to be used for the conduct of the proposed research should be briefly described. Indicate the pertinent capabilities of the institution, including support facilities (such as machine shops), that will be used during the project. List the most important equipment items already available for the project and their pertinent capabilities. Include this information for each subcontracting institution (if any).

### **3.10 Other Support of Investigators**

Other support is defined as all financial resources, whether Federal, non-Federal, commercial, or institutional, available in direct support of an individual's research endeavors. Information on active and pending other support is required for all senior personnel, including investigators at collaborating institutions to be funded by a subcontract. For each item of other support, give the organization or agency, inclusive dates of the project or proposed project, annual funding, and level of effort (months per year or percentage of the year) devoted to the project.

### **3.11 Appendix**

Information not easily accessible to a reviewer may be included in an appendix, but **do not use the appendix to circumvent the page limitations of the proposal**. Reviewers are not required to consider information in an appendix, and reviewers may not have time to read extensive appendix materials with the same care they would use with the proposal proper.

The appendix may contain the following items: up to five publications, manuscripts accepted for publication, abstracts, patents, or other printed materials directly relevant to this project, but not generally available to the scientific community; and letters from investigators at other institutions

stating their agreement to participate in the project (do not include letters of endorsement of the project).

#### **4. Detailed Instructions for the Budget**

(DOE Form 4620.1 "Budget Page" may be used).

##### **4.1 Salaries and Wages**

List the names of the principal investigator and other key personnel and the estimated number of person-months for which DOE funding is requested. Proposers should list the number of postdoctoral associates and other professional positions included in the proposal and indicate the number of full-time-equivalent (FTE) person-months and rate of pay (hourly, monthly or annually). For graduate and undergraduate students and all other personnel categories such as secretarial, clerical, technical, etc., show the total number of people needed in each job title and total salaries needed. Salaries requested must be consistent with the institution's regular practices. The budget explanation should define concisely the role of each position in the overall project.

##### **4.2 Equipment**

DOE defines equipment as "an item of tangible personal property that has a useful life of more than two years and an acquisition cost of \$25,000 or more." Special purpose equipment means equipment which is used only for research, scientific or other technical activities. Items of needed equipment should be individually listed by description and estimated cost, including tax, and adequately justified. Allowable items ordinarily will be limited to scientific equipment that is not already available for the conduct of the work. General purpose office equipment normally will not be considered eligible for support.

##### **4.3 Domestic Travel**

The type and extent of travel and its relation to the research should be specified. Funds may be requested for attendance at meetings and conferences, other travel associated with the work and subsistence. In order to qualify for support, attendance at meetings or conferences must enhance the investigator's capability to perform the research, plan extensions of it, or disseminate its results. Consultant's travel costs also may be requested.

##### **4.4 Foreign Travel**

Foreign travel is any travel outside Canada and the United States and its territories and possessions. Foreign travel may be approved only if it is directly related to project objectives.

##### **4.5 Other Direct Costs**

The budget should itemize other anticipated direct costs not included under the headings above, including materials and supplies, publication costs, computer services, and consultant services (which are discussed below). Other examples are: aircraft rental, space rental at research

establishments away from the institution, minor building alterations, service charges, and fabrication of equipment or systems not available off-the-shelf. Reference books and periodicals may be charged to the project only if they are specifically related to the research.

#### **a. Materials and Supplies**

The budget should indicate in general terms the type of required expendable materials and supplies with their estimated costs. The breakdown should be more detailed when the cost is substantial.

#### **b. Publication Costs/Page Charges**

The budget may request funds for the costs of preparing and publishing the results of research, including costs of reports, reprints page charges, or other journal costs (except costs for prior or early publication), and necessary illustrations.

#### **c. Consultant Services**

Anticipated consultant services should be justified and information furnished on each individual's expertise, primary organizational affiliation, daily compensation rate and number of days expected service. Consultant's travel costs should be listed separately under travel in the budget.

#### **d. Computer Services**

The cost of computer services, including computer-based retrieval of scientific and technical information, may be requested. A justification based on the established computer service rates should be included.

#### **e. Subcontracts**

Subcontracts should be listed so that they can be properly evaluated. There should be an anticipated cost and an explanation of that cost for each subcontract. The total amount of each subcontract should also appear as a budget item.

### **4.6 Indirect Costs**

Explain the basis for each overhead and indirect cost. Include the current rates.