

Frequently Asked Questions: DE-FOA-0001059, LAB 14-1059
Resilience for Extreme Scale Supercomputing Systems
Posted August 3, 2014
Revised August 13, 2014

This document will be updated with answers to the frequently asked questions that are received – please visit this page occasionally.

A number of questions address the technical scope of the Announcements. Given a relatively small budget in relation to the breadth of research challenges related to resilience, the scope of the Announcements is intentionally restricted to focus on foundational research to advance our understanding of the resilience challenges and the opportunities that hardware indicators might afford, in hopes of making significant progress in these areas. We recognize that additional approaches to resilience are potentially valuable and we hope that a later Announcement broader in scope can address other opportunities.

Q-1: The second question is about the Topic 2: Fault Mitigation. The FOA asks several questions:

- **"How can an application establish its tolerance for error and communicate that information to the system components handling errors?"**
- **"If applications can tolerate or handle soft errors, how can they best convey that to the appropriate layer(s) of the software stack? What mechanisms can be provided to help applications to better handle soft errors?"**
- **"How can we reduce the amount of code that is impacted by fault mitigation?"**

The Announcement also states:

"Out of scope for this FOA are applications that ... Aim to incorporate resilience into programming models and/or languages.

Q-1.1: Can you say more about the restriction on resilience for programming models and languages? We assume that means you are not interested in something like fault tolerant MPI, but most software-based solutions have to be implemented in some model. Where is the line?

A-1.1: Developing a domain-specific language for resilience, language enhancements for resilience, and/or other programming model approaches that address language expressiveness with respect to resilience are out of scope. Those important topics are addressed by other Announcements.

Q1.2: Are APIs for resilience (libraries and annotations) considered out-of-scope?

A-1.2: APIs and libraries for resilience are within scope.

Q-2: With respect to Topic 3: Anomaly Detection for Fault Avoidance, the Announcement says: "Research is needed to determine whether it is possible for machine learning methods to

enable anticipation of faults and/or node failure far enough in advance for an application to avoid the fault/failure, perhaps by migrating the running application to another node."

The Announcement also states that, "Out of scope for this FOA are applications that: ... Address fault prediction through machine learning other than by anomaly detection based on information from hardware sensors"

Q. 2.1: Can one include (in addition to hardware sensors) the use or development of software sensors to track and detect higher level phenomena, like for example application performance variations, congestion in file system or other macro behaviors not tractable at hardware level?

A-2.1: No, the use or development of software sensors to track and detect higher level phenomena, like for example application performance variations, congestion in file system or other macro behaviors not tractable at hardware level are out of scope.

Q-2.2: We believe that an integrated fault prediction system should use any possible source of information -- whether it comes from hardware or software -- in particular, any information that is collected in current error logs. The current constraint will arbitrarily restrict the sources of information.

A-2.2: See A-2.1 above. The constraints are not arbitrary; they reflect ASCR priorities within the funding available. Developing an integrated fault prediction system is not a goal for this Announcement. We hope to address this challenge with a later Announcement.

Q-2.3: Does evaluating data from hardware performance counters qualify as using "hardware sensors" for theme 3?

A-2.3: Yes.

Q-3: The FOA says that "Requests for supplements to existing projects to address resiliency challenges" are out of scope. What about resilience projects that collaborate with application teams to make sure their work addresses real application use-cases and includes some funding for the application teams? Are such projects completely out of scope or do we need to ensure that such projects do not develop resilience mechanisms for the individual applications but rather more general techniques.

A-3: The intent is twofold:

- Existing projects may not apply for supplements to the existing awards to enhance resilience of their applications. The 14-1059 Announcements focus is on developing new methods to achieve resilience.

- The techniques developed under this Announcement should be generalizable to more than one application. Collaborating with applications to ensure that realistic use-cases drive the work is encouraged.

Q-4: To be application-agnostic, would work with 3+ codes be sufficient (in different science areas and different solvers)?

A-4: Working with three applications in different science areas is sufficient. Research aimed at developing resilient solvers and/or math libraries is out of scope, having been addressed by another recent call.

Q-5: What are the main deliverables? Is this a research-based effort where ASCR would like to see reports and papers, or would it be better to have a delivered software-package?

A-5: The deliverables should be appropriate to the proposed research. For some lines of research, reports and papers might suffice. For others, it is hard to imagine substantive progress without development of at least prototype software.

Q-6: Since this is exascale-focused, is working with applications to tests the proposed methods of research which can scale today to 10-100k cores sufficient? (Where's the break point?)

A-6: Since we don't have an exascale platform yet nor applications that would run efficiently on such a machine, applicants should describe the applications and platforms they intend to use in the research. The "break point" is feasibility for the timeline of this Announcement.

Q-7: Where can I find more information about the supercomputing facilities that are supported by ASCR?

A-7: General information about ASCR-supported facilities is available here, including links to information about each facility and the relevant allocation policies:
<http://science.energy.gov/ascr/facilities/>.

We have received a variety of questions regarding the composition of collaborative teams:

Q-8: Is there a benefit or disadvantage to working with an industry partner? (E.g., does ASCR feel it is already funding industry research into resilience through FF and DF contracts?) Can companies (Intel, Nvidia, ISI, etc.) and/or non-profits (HDF-group, etc.) participate? Can they receive funding (through sub-contracts)?

A-8: Working with an industry partner is encouraged if doing so benefits the proposed research but it is not required.

Not-for-profit organizations may also apply. Both industry and not-for-profit organizations can apply directly to DE-FOA-0001059 and are not required to be sub-awardees. They can also serve as Lead institutions.

Q-9: Is a proposal involving two universities but not a DOE National Laboratory still regarded as a collaborative proposal? Does it make any difference whether there are separate proposals from each university or sub-awards to a single university?

A-9: Yes, this would be a collaborative proposal in either case. DOE prefers to make separate awards to each participating institution. However, only the lead needs to submit a pre-application.

Q-10: Since some University collaborations will be smaller than \$100k (a single GSRA) can this be handled through a subcontract? Or is it ok to have a University collaborator in at this level?

A-10: ASCR prefers to have a higher level of commitment from a University than a single graduate student. If the single graduate student is a student of a DOE National Laboratory applicant with a dual appointment, the student should be supported via the Laboratory award.

Note that DOE does not constrain University faculty applicants to summer support only nor to only two weeks of support during the summer.

A few questions pertain to constraints on the number of applications/proposals from an institution or individual

Q-11: Is there a limit on the number of pre-proposals/pre-applications that an institution may submit, either as the lead or as a collaborator?

A-11: No.

Q-12: Can an individual participate in more than two pre-proposals/pre-applications?

A-12: Yes. This is to allow for the possibility that one or more submissions might not result in encouragement to provide a full proposal/application. Choices would have to be made by the applicant if more than two pre-proposals/pre-applications were to be invited; that decision will not be made by DOE.

Eligibility of various organizations and the roles they can play have been the subject of several questions:

Q-13: Can a university be the lead institution on a proposal that includes a DOE National Laboratory as a collaborator

A-13: Yes.

Q-14: Can an industry applicant be the lead for a proposal?

A-14: Yes.

Q-15: Is a Federally Funded Research and Development Center (FFRDC) with primary sponsorship by a Federal agency other than DOE eligible to apply as a sub-awardee to a DOE National Laboratory?

A-15: Yes.

Q-16: Is a non-USA institution eligible to apply as a Lead or collaborating organization?

A-16: Yes. However, a non-domestic applicant, to be able to compete, should include a demonstration of unique skills, abilities, and resources that are not available in domestic applicants. In general, we advise non-US organizations that their greatest likelihood of success is being a sub-awardee on projects proposed and led by DOE National Laboratories.

Q-17: Is it a requirement that an organization other than a DOE National Laboratory must partner with a DOE National Laboratory?

A-17: No. Partnerships should be based on benefit to the proposed research.

Issues related to submission of pre-applications/pre-proposals are addressed here:

Q-18: Is the size of the pre-application abstract the usual 2 pages maximum, including references?

A-18: Yes, the application abstract should not exceed 2 pages, including figures and references. However, the limit does not include the conflict-of-interest list.

Q-19: If a pre-proposal/pre-application results in an invitation to submit a full proposal, can the budget be changed?

A-19: Yes.

Q-20: If a pre-proposal/pre-application results in an invitation to submit a full proposal, can the list of participants be changed?

A-20: The full proposal cannot include additional co-investigators or senior personnel who contribute intellectual content to the proposal, given our use of the pre-proposal/pre-application COI list to begin developing lists of reviewers. However, the proposal may indicate the intent to support an unidentified senior person to provide particular expertise to the project. Note that junior personnel such as programmers can be identified later and may not be identified by the time of proposal submission.

Q-21: If a collaborative pre-proposal/pre-application results in encouragement to submit a full proposal, can the Lead PI/organization be changed among the participants in the pre-proposal/pre-application?

A-21: No. The Lead PI/organization for the encouraged full proposal must match those on the submitted pre-proposal/pre-application.