

Request for Information Regarding HPC Procurement at NSC

2017-06-16

Introduction

This Request For Information (RFI) pertains to upcoming procurements of high performance computing (HPC) resources AC2018 and NWP18 for National Supercomputing Centre (NSC) at Linköping University, Sweden. The procured systems will be delivered and put in operation during 2018. The deadline for answering this RFI is **2017-07-04**. It is not mandatory to reply to this RFI but we would appreciate your views and comments regarding these procurements.

Your answers to this RFI will be treated as “working material” and are confidential until contracts have been awarded.

To put this RFI in some context, a high level description of the procurements is outlined in this document. However, some of the below questions are directly related to details not found in this RFI. For these details about the procurements, see the procurement pre-announcements for AC2018 and NWP18 published alongside this RFI. These pre-announcements can be found at **TED**, www.opic.com and the material is available at www.nsc.liu.se/about/procurement.

AC2018 and NWP18 are two separate procurements, but due to the similarities between them, this RFI relates to both procurements. In case your answer differs between the two procurements, please indicate which procurement your answer is valid for.

For these procurements, the intention is to stay agnostic with respect to what hardware delivers the computational power, specifically with respect to what CPU architecture is used, but also with respect to what interconnect is used, and in the case of AC2018, GPU accelerators are a vendor option to a certain fraction of the system budget. The evaluation of offers will be based on job throughput over TCO.

The purpose of this RFI is to solicit background information from interested vendors and subcontractors, thereby permitting better requirements in technical specifications as well as in vendor qualification.

AC2018 Description

The HPC resource AC2018 will be a general purpose HPC resource with a throughput focus, targeting academic research in predominantly computational materials physics and computational chemistry.

Anticipated hardware:

- CPU Architecture
 - x86_64 (AMD, Intel)
 - POWER
- GPU capabilities
 - CUDA capable
 - ECC memory
 - Paging to main memory
- Interconnect
 - InfiniBand
 - Omnipath

NWP18 Description

The NWP18 HPC resource will consist of two installations and will run ensemble weather forecasts for the Norwegian, Swedish and Finnish meteorological offices within MetCoOp.

Anticipated hardware:

- CPU Architecture
 - x86_64 (AMD, Intel)
 - POWER
- Interconnect
 - InfiniBand
 - Omnipath

Contact Information

Contact person for the procurements of AC2018 and NWP18:

Name: Katja Ekström
Position: Procurement officer
E-mail: katja.ekstrom@liu.se

Questions

1. Who are you?
 - a. Company name
 - b. Contact person for further discussion regarding your answers to the following questions.
 - c. Annual revenue of the part of your company that is related to the business of HPC
2. What solutions have you provided as a contract partner:
 - a. Which HPC systems have you delivered in Europe during the last three (3) years (more than five systems need not be listed)?
 - b. Have you delivered any system(s) entering the Top500 list the last three (3) year and on which position(s) in such case (more than five systems need not be listed)?
 - c. Have you delivered any HPC system(s) in Sweden and of what size(es) in such case?
3. Which NSC procurements do you intend to participate in (more than one can apply):
 - a. AC2018
 - b. NWP18
 - c. None
4. In which role do you intend to participate? (more than one can apply):
 - a. contract partner
 - b. subcontractor
 - c. manufacturer
 - d. integrator
 - e. reseller
 - f. service provider and/or logistics of spare-parts
 - g. support
 - h. other, please specify

If you are going to tender:

5. Describe your intended partner structure and their responsibilities with regards to the following activities:
 - a. benchmarking
 - b. manufacturing
 - c. integration
 - d. hardware installation
 - e. functional test and accept
 - f. service
 - g. support

- h. spare part logistics
6. Which CPU-architectures can you provide in a contract?
 - a. Intel Xeon
 - b. Intel Xeon Phi
 - c. AMD Zen
 - d. POWER
 - e. ARM
 7. Which cooling solutions can you provide in a contract?
 - a. Air cooled
 - b. Water cooled racks
 - c. Water cooled servers
 - d. Immersed
 - e. Other, please specify
 8. Which range of installed densities of servers and processors per rack do you provide for the different CPU-architectures and cooling solutions? (1 rack is assumed to be 600 mm wide and 2000 mm tall (42U) unless stated differently in your answer).
 9. What time to installation can you offer from the time of contract signing?
 10. What degree of integration and system QA testing is performed in the factory before shipping?
 11. What type and level of service and support would you propose in a contract of the size of these procurements?
 - First and foremost we want to keep the level of faults as low as possible. Installed equipment should provide its service.
 - To reduce cost we want to reduce and limit the time personnel at NSC spend on handling typical, recurring errors and failures (e.g. memory faults, PSUs) during the life-time of the system. This includes not only replacing the component but also time spent on fault diagnostics, error reporting and ticket handling, logistics of parts (RMA handling), waiting for response, etc. E.g. replacing a broken dimm from a on-site cache of spare parts may be more efficient than having service personnel coming on site. But, the devil is in the details here, how each task is managed, if it is fast or time-consuming for NSC personnel.
 - We also need to have some protection if the fault frequency becomes higher than normal, temporarily or constantly. The vendor should take responsibility for the additional cost delivered equipment can incur if systematic errors occurs. E.g. high failure rate of components due to quality issues in factory, servers that are unable to pin-point faulty memory DIMM, general compatibility issues, etc.
 - Furthermore we need an efficient escalation management of systematic errors to resolve these issues.

The required service level can be divided into two categories:

1. Critical equipment such as central equipment which is either rather unique and/or essential for the delivery of the computing services.
2. The sheer bulk of computing capacity.

The first category may require spares on-site to resolve issues swiftly. For the bulk of computing it may be more efficient to increase the bulk of servers with a few to facilitate a more efficient scheduling of fixes, replacement activities and shipping of parts instead of handling each issue individually.

13. Other comments?

Request for Comments and Suggestions

Co-located with this RFI, there is an early draft of the invitation to tender. This is by no means finished at this point in time. If you have comments and/or suggestions for improvements to make this procurement more transparent, unambiguous and less intricate, you are very welcome to contribute. We will consider your comments and suggestions when finalizing the specification and requirements for the procurement.