PDC Center for
High Performance Computing

Stockholm, October 25, 2009

Invitation to submit proposal for a new HPC resource at PDC/KTH (within "ramavtal")

It is our pleasure to invite you to submit a proposal for a new HPC resource to be installed at PDC/KTH in Stockholm.

Written proposals marked CFP200910 on the envelope should be delivered to Erwin Laure at the address below, no later than 24:00, January 5, 2010. As specified in the information attached, the evaluation of proposals will be based on both expected performance and economical factors, most importantly the expected computational capacity, energy efficiency, and time to production usage. A proposal can be accepted without prior negotiation.

We expect to finalize the process in January 2010 and would prefer to have the physical installation finished before the end of Q1 or early Q2, 2010. The resource should be available for production usage by the end of June 2010.

Requests for further information can be addressed to Erwin Laure erwinl@pdc.kth.se (+46-8-790 65 14) or Daniel Ahlin, dah@pdc.kth.se (+46-8-790 68 44) for technical questions.

Erwin Laure
Director PDC-HCP
KTH

1. Background

PDC is one of the major Swedish supercomputing centers and provides academic researchers with world-class compute and storage facilities. PDC contributes to the Swedish National Infrastructure for Computing (SNIC) and this procurement is funded by SNIC with the aim to make a general-purpose HPC resource available to SNIC users. The system is intended for users requiring a large degree of parallelism and will be also used as application evaluation system for a potential future Swedish system in the multi-100 teraflops range. The budget for this procurement is approx. 20.5 MSEK.

2. Overview

The new system should be a general purpose HPC system for usage by a wide variety of scientific users, including codes from molecular dynamics, CFD, computational chemistry, and bioinformatics. The system should be able to run codes optimized for x86_64 architectures efficiently. The node/network performance should be balanced to allow wide jobs covering the whole machine running efficiently. Particularly, it should support MPI jobs with at least up to 10,000 processes with at least 1.5 - 2 GB RAM/process.

In addition to the HPC machine itself we request offers for matching cooling solutions (in-row, back-door, liquid or similar).

The budget of approx. 20.5 MSEK is supposed to cover both the computing system and cooling.

3. Requirements

The requirements are:

1. The system should provide high performance for parallel scientific applications optimized for x86_64 architectures using up to the entire machine. The following benchmarks will be used to assess the capabilities of the system: stream, linpack, NAS parallel benchmarks, Intel MPI benchmarks. We will also use application codes like VASP (http://cms.mpi.univie.ac.at/vasp/), Dalton (http://www.theochem.kth.se/dalton/), Gromacs (http://www.gromacs.org/), and SIMSON (http://www.mech.kth.se/mech/info_tritamek.jsp?TritaMekID=315) for benchmarking.

2. The system should be energy efficient. We regard power efficiency equally important to the overall performance of the system. Energy efficient versions of CPUs should be considered as well as diskless solutions.

3. The communication network shall be high performance and balanced with the node performance. It should deliver high application performance when using MPI. We expect the network to deliver at least a bandwidth of 2,700 MB/s and a latency of no more than 2μs between any two nodes when using MPI or demonstrate performance equivalent to a network with these characteristics using the above mentioned application benchmarks.

4. Any software (e.g. OS, MPI) required to run the system has to be included or freely available. In the latter case the system should be compatible to recent Linux distributions.

KTHPDC-HPC Teknikringen 14, SE-100 44 Stockholm. Phone: +46 8-790 60 00. Fax: +46 8-24 77 84.
Pg No: 1 56 53-9. Bg No: 895-9223. VAT registration No: SE202100305401.

5. Appropriate data and management networks need to be available.

6. The cooling solution should be neutral to ambient cooling and integrated in PDC's existing cooling solutions. PDC is in the process to implement a heat reuse system that would require water temperatures of 13°C incoming and 23°C outgoing.

7. The system shall be in full production usage as of July 1st 2010. Running Linpack for the Top500 list will be one of the proofs of production readiness.

4. Information Required

For the computing system and cooling we need

1. A general description of the hardware, and possibly motivations for the choice if you believe an alternative solution would fit our needs.

2. A detailed list with hardware (and possibly software) specifications.

3. Costs including 3 years of service. The cost should be detailed for individual components (such as computing hardware, network, cooling, software, service).

4. Expected timetable for delivery and provided support for reaching production quality.

5. State the operating requirements for the proposed system: delivery, assembling, power consumption, cooling, size, service space requirements, weight, etc.

5. Selection Process

All proposals confirming to the requirements stated in Section 3 will enter the selection process. Important evaluation criteria will be (in no particular order or weighing)

1. Overall system performance

2. Energy efficiency

3. Readiness to offer guarantees of entering production by July 1st, 2010

4. Total cost of ownership (including e.g. space, electricity, cooling, service, administration effort)

Item 1) will be assessed using the benchmarks mentioned in Section 3 item 1) and item 2) by using High Performance Linpack as benchmark.

KTHPDC-HPC Teknikringen 14, SE-100 44 Stockholm. Phone: +46 8-790 60 00. Fax: +46 8-24 77 84.
Pg No: 1 56 53-9. Bg No: 895-9223. VAT registration No: SE202100305401.

6. Acceptance tests

The acceptance test is simplified and will be fully specified in the order. Roughly the acceptance test will be to demonstrate that all parts of the system operate according to the requirements. This includes a multi-user test to demonstrate that the system can handle many simultaneous users, a linpack run to list the machine on the top500 list, and the compliance to negotiated thresholds for application performance and energy consumption.

KTHPDC-HPC Teknikringen 14, SE-100 44 Stockholm. Phone: +46 8-790 60 00. Fax: +46 8-24 77 84.
Pg No: 1 56 53-9. Bg No: 895-9223. VAT registration No: SE202100305401.