

Proposal for a Danish National HPC Facility

18 June 2013

This is a joint proposal by:

Aarhus University (AU) and the University of Southern Denmark (SDU)

Executive Summary

Our vision is to establish a new national HPC facility, jointly funded by DelC and the participating universities. All other danish institutions are invited to join this proposal. The initial setup cost for the new Tier1 machine is fixed to 35M dkk, which includes the funds available through DelC. The University of Southern Denmark is committed to provide a physical location for the national facility, should it be placed within national borders.

Initial funding and construction

Adding to the funding of 15M dkk available through DelC, the participating institutions will invest 20M dkk into a Tier1 installation, which will form the core of the new national HPC facility. **The total cost for the envisioned Tier1 machine is fixed to 35M dkk.** This cost is deemed adequate for the provision of a Petaflop scale machine, competitive at the European and global level.

Adequate storage capacity to guarantee a seamless operation of the new machine is also considered part of the new installation. Network connectivity will be provided by the host institution. It is advised that the Tier1 machine will be deployed in two or possibly more phases, as common practice for large Tier1 installations.

At present, two universities are part in this proposal: the University of Southern Denmark (SDU) and Aarhus University (AU). SDU will invest 15M dkk in the initial setup costs, and AU will contribute with 5M dkk. **More institutions can and are welcome to join this partnership** to share the initial costs of the national facility. However, as mentioned above, the budget for the Tier1 machine will be kept fixed to 35M dkk. The dean of the faculty of engineering and science of Aalborg University, Eskild H. Nielsen, has expressed interest in the proposed national facility and AAU is resolved to participate in the forthcoming discussion for the establishment of the Tier1 centre.

We strongly favor a single national HPC solution, given the insufficient level of funding for multiple installations and the fact that a single Tier1 machine is surely enough to cover the needs of Danish users.

It is understood that the institutions, which have initially invested in the setup costs of the national centre, will benefit from a free quota of computer time on the Tier1 machine during the first 4 years of operations of the new HPC centre. The amount of the free share is based on the initial investment of each institution, as described below in this document. In total about 1/4 of the machine will be reserved for the initial investors during the first 4 years of operations.

The specific type of hardware to be procured must be agreed upon by the Change Advisory Board (CAB), composed by HPC users and representatives from the research institutions. We suggest below some general guidelines for the hardware configuration of the Tier1 machine.

The CAB should also assess the profitability, with respect to running costs, of locating the new HPC facility in a foreign country as opposed to Denmark. If placed in Denmark, the parties in this proposal agree that the new facility will be hosted at the University of Southern Denmark (SDU). SDU can provide the space and appropriate infrastructure either in a building on the main university campus in Odense or via the use of dedicated containers, subject to the opinion of the CAB.

Operation model

Access to the new national HPC centre, both from the academia and enterprises, will be charged a fee per node*hour. In agreement with the guidelines from DelC, ~60% of the users' charges will cover the running costs, which include electricity, cooling and network, while the remaining ~40% will be used for the renovation of the HPC infrastructure. Estimating that the annual running expenses are, at most, about 1/3 of the cost of the machine, in the proposed financing model a complete renovation of the hardware is possible every ~5 years, at 80% utilization of the machine. An upgrade of the computing hardware every 4-5 years is in line with the current practice for updating the major European Tier1 sites, such as the Jülich Supercomputer Centre (Germany), CINECA (Italy), Barcelona Supercomputing Centre (Spain) or the centers run by CEA (France), just to name a few. The details of the renovation procedure should be decided by the CAB.

The investments for the initial construction of the machine will be taken into account by granting free computer time to the initial investors of the Tier1 machine. The amortization plan for the initial costs will have a duration of 4 years, as described in more detailed below (see section "*Illustration of the proposed operation model*").

We also propose that the new national Tier1 machine become part of PRACE, with modalities to be decided by the CAB.

The possibilities offered by the national Tier1 installation should be advertised to small and medium enterprises and its exploitation throughout the Danish industry should be encouraged.

Given its strategical value for the whole country and its potential to accelerate the innovative process, we also urge all the interested parties to seek a long-term co-financing plan for the national HPC facility at the national level. We also remark that the lack of a common national strategy in the past results in Denmark being today the only nordic country without an adequate national HPC infrastructure, such as the proposed Tier1 site.

Guidelines for the Tier1 system hardware

The CAB will be responsible for the choice of the specific type of hardware to be procured for the Tier1 system. In this section we suggest a few general guidelines for the Tier1 system configuration.

It is crucial that the new Tier1 system is able to accommodate as many scientific and industrial projects as possible. However it is also of the utmost importance to maintain the hardware configuration of the machine as uniform as possible. Different kind of accelerators, such as GPUs or Intel Phi, can easily be accommodated within a common modular solution, readily available through several hardware providers. It is also critical that the computing nodes have fast interconnections.

Adequate scratch disk space must be available to users, to ensure the adequate operations of the Tier1 centre. However it is understood that the new national Tier1 facility is not a data storage facility. Also fast internet connectivity is needed to move the data in and out of the Tier1 facility.

For purely illustrative purposes, a possible configuration for a Tier1 machine is the following: 1000 computing nodes (16 cores, 64GB memory), ca. 22 racks; 1/3 of the nodes equipped with 2 GPU accelerators (e.g. NVidia K20) and 1/3 with 2 CPU accelerators (e.g. Intel Xeon Phi); 150TB local SSD scratch space; 40Gbps Infiniband Interconnect; 560TB of storage scratch space, ca. 1 rack. The total computational speed (peak performance) of this configuration is about 1.8 PFlops (double precision). The estimated energy consumption is about 850 KW (including cooling).

Illustration of the proposed operation model

The goal of this section is to present, in a very simplified but transparent way, how the proposed operation model should work. While the figures used in this section are believed to be reasonable estimates of the costs relative to a Tier1 facility located in Denmark, they are only provided for illustrative purposes at this stage. Actual costs depend on the particular choice of the hardware and must be assessed in detail by the CAB.

Here we assume that the total annual running costs for the machine are 1/3 of the hardware cost, indicated below by X . With the aim of refinancing the machine every 4 years at 100% capacity, the total cost of ownership for the machine for 1 year is:

$$\text{annual cost} = \text{running cost} + \frac{X}{4} = \frac{X}{3} + \frac{X}{4} = \frac{7}{12}X$$

At 80% capacity, the machine can be refinanced in 5 years, as mentioned above. For a 35M dkk machine the total ownership cost for one year is then ~20M dkk, which corresponds to a cost of about 2.3k dkk per hour. The latter is thus the corresponding user fee for the *complete* machine for 1 hour. Given that the Tier1 machine will have a number of the order of 1k nodes, the user fee amounts to about 2.3 dkk per node*hour.

The initial investments of the participating institutions can be amortized over the first 4 years of operations by reserving a fraction of computer time on the Tier1 machine for these institutions free of charge. If we indicate with Y the total initial investment of the participating institutions, then the fraction of the machine reserved for the investors during the first 4 years is

$$\text{fraction reserved for initial investors} = \frac{Y}{4} : \frac{7}{12}X = \frac{3}{7} \frac{Y}{X}$$

Based on the total investment of the participating institutions, then the fraction of the machine reserved for the investors during the first 4 years free of charge is about 25% of the machine. This quota of computer time will be split among all initial partners proportionally to their investment.

The amortization plan also implies that, during the first 4 years, the users from other institutions will pay more than the previously calculated fee, to keep the total refinancing time fixed as described above. Using the figures above the user fee will be increased by about 33% to a total of about 3 dkk per node*hour during the first 4 years.

Summary

With this proposal we suggest a Tier1 installation that may serve as a national HPC facility which is setup and financed in accordance with the guidelines given by DelC.

Henrik Pedersen

Dean, professor, Faculty of Natural Sciences, University of Southern Denmark

Brian Bech Nielsen

Dean and forthcoming rector, professor, Science and Technology, Aarhus University

This proposal is supported by the following scientists at AU and SDU:

Aarhus University:

Assoc. prof. Ove Christiansen, Dep. of Chemistry

Assoc. prof. Frank Jensen, Dep. of Chemistry

Prof. Poul Jørgensen, Dep. of Chemistry

Prof. Poul Nissen, Dep. of Molecular Biology and Genetics

Prof. Jeppe Olsen, Dep. of Chemistry

Prof. MSO Birgit Schiøtt, Dep. of Chemistry

University of Southern Denmark:

Assoc. Prof. Jan Baumbach, Dep. of Mathematics and Computer Science

Assoc. Prof. Michele Della Morte, Dep. of Mathematics and Computer Science

Assoc. Prof. Hans J. A. Jensen, Dep. of Physics, Chemistry and Pharmacy

Prof. Ole Nørregaard Jensen, Dep. of Biochemistry and Molecular Biology

Assoc. Prof. Himanshu Khandelia, Dep. of Physics, Chemistry and Pharmacy

Assoc. Prof. Jacob Kongsted, Dep. of Physics, Chemistry and Pharmacy

Prof. Ole G. Mouritsen, Dep. of Physics, Chemistry and Pharmacy

Prof. MSO Claudio Pica, Dep. of Mathematics and Computer Science

Prof. Steen Rasmussen, Dep. of Physics, Chemistry and Pharmacy

Prof. Francesco Sannino, Dep. of Physics, Chemistry and Pharmacy

Prof. Achim Schroll, Dep. of Mathematics and Computer Science

Assoc. Prof. Iliia A. Solov'yov, Dep. of Physics, Chemistry and Pharmacy

Assoc. Prof. Carsten Svaneborg, Dep. of Physics, Chemistry and Pharmacy

Assoc. Prof. Fabio Vandin, Dep. of Mathematics and Computer Science