

Evaluation of computing needs at Linnaeus University

Monday 4th May, 2015

1 Introduction

This document shall summarise the needs for computing resources of the different departments and groups at Linnaeus University.

In order to collect information from the groups, a short, informal workshop was organised by Yvonne Becherini on Friday 24th April, 2015. The workshop agenda as well as some of the presented talks can be found at <https://indico.cern.ch/event/388005/>

1.1 Participants

- Yvonne Becherini (Organizer, Chair)
- Andreas Haupt (Astroparticle Physics) - notes
- Magnus Paulsson (Condensed Matter Physics)
- Sabri Pllana (Computer Science)
- Aris Alissandrakis (Media Technology)
- Andreas Linderholt (Mechanical Engineering)

Due to the short-term organisation of the meeting some people could not personally attend but kindly provided slides:

- Alf Månsson (Bionanotechnology)
- Ian Nicholls (BBCL)
- Jarone Pinhassi (EEMiS)

2 Computing demands of different groups at LnU

2.1 Astroparticle Physics

The main large-scale computing tasks of the Astroparticle-physics group are Monte-Carlo simulations (e.g. Air showers of cosmic particles, studies of Cherenkov-telescope electronics) as well as data reduction and analysis for the H.E.S.S. and Fermi experiments. Whereas simulations are clearly cpu-bound, i.e. the demands for storage are rather low, data analysis tasks demand a fast, reliable file system. All in all 100 modern cpu cores with 4GB ram per core (provided by a batch queueing system) and 100-500TB of fast mass-storage should fulfil the computing demands of these tasks. As the data rates (both: reads

and writes) need to scale with the amount of simultaneously running jobs, it is strongly required to base the storage solution on a parallel file system like "GlusterFS", "Lustre" or another adequate product.

The used software is developed within the collaborations and its code quality varies rather heavily. As the preferred operating system in other computing centres supporting the collaborations is RHEL-based, running on a compatible system is highly preferable. Software tests are typically only done on these Linux distributions. In the current state the software is only single-threaded and does not provide modern methods of parallelisation. However, multiple, independent tasks can be run simultaneously.

In order to develop and test the collaboration software, the group requires a so-called WGS ("Working-Group Server") which is set up identical to the execution nodes provided in the batch system. This WGS must provide (external) logins via ssh.

Additional to the large data storage, space for installing software group-wide is required ("group space"). This one needs to be visible on all nodes (WGS, batch nodes) and should provide some 100s of gigabytes of storage. Furthermore personal space for each group member in the order of a few gigabyte (\$HOME directory) is required.

In order to provide compute resources to the new CTA ("Cherenkov Telescope Array") collaboration, grid-based access to the batch system is mandatory as e.g. simulation campaigns are organised and run centrally.

All requirements listed up to this point can be summarised under the use-case "large-scale computing". The SNIC computing centre in Lund (Lunarc) currently provides these resources. However, due to bad luck there was a major data loss of Lunarc's mass-storage system last year connected with a longer outage of the service. The provided SNIC storage ("SweStore") does not fit the use-case of a fast parallel file system in its current state, unfortunately. It is mainly intended for data archival.

Additional to the large-scale computing demands, the group needs a Wiki for group-internal documentation as well as access to a meeting-software like "Indico". Personal web space to share e.g. automatically generated histograms, plots, etc. with the collaboration is currently missing, as well. These services are not provided by SNIC.

2.2 Condensed Matter Physics

The group consists of 4 people and uses compute resources mainly for electronic structure calculations. The used software is Linux-based and parallelised using MPI. It requires a low-latency network (e.g. Infiniband) if run on multiple nodes. Memory requirements are rather low: ≈ 2 GB per core. Additionally, each user requires 200GB of storage.

In order to run tests and code development, a small Linux cluster named "Bionano" was set up in the Kalmar computing centre five years ago, funded by "Sparbanksstiftelsen". It consists of 4 eight-core "Intel Xeon" nodes with ≈ 1 TB local data which are backed up on a daily basis. The cluster provides shared access (for the Physics, Biology and Chemistry groups) and is officially maintained by Magnus. As these nodes run without problems up to now and the installed Linux distribution is no longer maintained, this does not result in much work currently. Nevertheless, the nodes are quite old and hardware failures (especially the hard disks) are to be expected in the next future.

The test runs are done rather seldom, so there is just a peak usage of 8-32 cpu cores for ≈ 1 day following a period without any usage for some months. The SNIC centre in Lund is then used for large-scale productions, again with only short times of activity. Medium-sized SNIC allocations (80000 cpu-core hours per month) are granted without problems. The group feels happy with the current solution.

2.3 Computer Science

The Computer Science group demands a computing facility in order to fulfil its teaching and research mandate. As an example software engineering in "Big Data" requires access to a compute cluster.

The proposed facility, called "LnU HPC Centre", should consist of a small Linux-based compute cluster accumulating ≈ 1000 cpu cores and 1TB of RAM. The idea would be to run it as shared resource, so other groups at LnU could use it for their purposes, as well. Additionally, a dedicated node for doing interactive visualisation (hosting: 2 recent cpus, 256-512GB ram, 8TB secondary storage, preferably on SSD) is required.

2.4 Media Technology

The group's computing demands involve a rather low volume of number crunching tasks for both, teaching and research. Currently most of these tasks are done using "Amazon Computing Cloud" services. Storage demands do not exceed the amount of 1TB, an increase is not expected.

2.5 Bionanotechnology

The computing needs of the Bionanotechnology group concentrate on simulations, image analysis and solving differential equations. These are done with the software products "Matlab" and "Femlab" on a workstation running a Windows operating system. Additional resources are not needed, currently.

2.6 Bioorganic and Biophysical Chemistry Laboratory (BBCL)

The BBCL works with molecular recognition phenomena and the informed design, development and application of functional and biomimetic materials. Areas of application include catalysis sensing and biomaterials. The current and projected computational activities include:

- Molecular Dynamics (primarily for complex mixtures)
- Quantum Chemistry (primarily for small molecule systems)
- Chemometrics/multivariate statistics (study of structure-function relationships using a Windows workstation)

In order to fulfil their large-scale computing demands, BBCL is using SNIC medium-size allocations ($\approx 50k$ cpu-core hours per year). Software being used include "Gaussian09", "AMBER" and "GROMACS".

The group especially demands space for storing larger amounts of data. Currently 50TB are needed, the yearly increase is expected to be 5-10TB.

2.7 Centre for Ecology and Evolution in Microbial model Systems (EEMiS)

The EEMiS group concentrates on bioinformatics analyses on genes (genomics), gene transcripts (transcriptomics) and proteins (proteomics) on model microbe species and microorganism assemblages in the natural environment (e.g. sea, sediments, soil, birds and fish).

For their analysis they make currently use of a Linux-based compute cluster at the "National Centre for Molecular Biosciences" (SciLifeLab) in Stockholm. The group's accumulated monthly usage varies between 10k-90k cpu-core hours. Single-node parallelisation (shared memory) is implemented, an increase of MPI usage is foreseen.

Data storage is a big issue, currently. Up to now 25TB of scientific data have been accumulated but a yearly increase of 40TB is expected. Long-term storage for data archival is a major issue. The allocated space at SweStore recently ran full.

Additional to their storage shortage, the group demands a web server front-end system for their databases, preferably running Linux.

2.8 Mechanical Engineering

The Mechanical Engineering group uses two workstations running under Windows 2008/R2 mainly for doing CAD ("Computer-Aided Design") tasks. Access is provided remotely via "Rdesktop". One of these nodes is intended to supply students access to a login node with all needed software installed. They are available only on Microsoft Windows operating systems and include:

- Abacus
- Nastran, Patran, Simexpert (vendor: "MSC software")
- Matlab
- Microsoft Excel

Additionally, the workstations provide ≈ 12 TB of data storage where currently $\frac{1}{4}$ is being actively used. There are no additional computing demands from this group.

3 Workshop discussion

Sabri Pillana presented his idea of establishing a shared compute cluster at LnU. It is intended to fill the gap between personal "notebooks" on one side, and "super computer facilities" provided by SNIC on the other side. Sabri emphasized, such an installation should become a service of whole LnU, rather than of one department. Some small and more or less private installations do even exist as of today. Nevertheless, other groups typically do not have any knowledge about them. There is a lack of dedicated system administrators from the IT department, additionally.

Especially participants from the Astroparticle and Computer Science group expressed their displeasure with the current situation. The IT department is unable to provide some required services. One of the identified issues is IT staff being experienced in administrating Windows-based systems, only.

It was questioned whether it is politically a good idea to finance SNIC for providing any larger computing resources on long-term. There are concerns, members of "non-SNIC universities" might even experience discrimination when applying for large resources. The SNIC board consists of members coming from "SNIC universities", only.

Investment costs for an initial installation should not exceed 2 million SEK. It was discussed whether at least parts of it could be financed by external funding. A possible co-financing with external institutes or industry should be investigated. Manpower is the critical aspect of this project. Whereas the purchase of the installation would be a one-time investment mainly, keeping experienced personnel at LnU is a permanent effort. It was agreed that once the system is fully set up, its administration will not exceed $\frac{1}{2}$ FTE. However, there are concerns whether it is a good idea to rely on a single person maintaining the installation. Furthermore, the Astroparticle group demanded large, reliable storage resources. Its maintenance is typically more manpower-intensive than running a compute facility, only.

It was decided to investigate the possibility of installing and running a medium-sized compute facility at LnU together with the LnU administration. According to initial investigations concerning the manpower costs of a full system administrator position, ≈ 800 k SEK are to be calculated per year.

4 Summary

This table briefly summarises the computing demands of LnU research groups:

Research Group	Task / Application Name	System Requirements
Astroparticle Physics	H.E.S.S./Fermi data analysis Simulations	Compute cluster, fast bulk storage Compute cluster
Condensed Matter Physics	Electronic structure calculations	Small compute cluster, MPI-capable
Computer Science	Software engineering	Compute cluster, dedicated visualisation node
Media Technology	Video generation	Small cluster or "Cloud computing", 1TB storage
Bionanotechnology	Simulations, Image analysis	Workstation with "Matlab" and "Femlab"
BBCL	Molecular Dynamics, Quantum Chemistry Chemometrics/multivariate statistics	Compute cluster, 50TB storage Windows workstation
EEMiS	Bioinformatics analyses	Compute cluster, 100TB long-term storage, Web server database front- end system
Mechanical Engineering	CAD	Windows workstation with CAD software