High Performance Computing Resource Allocations for Research Projects and External Partners

OVERVIEW AND GUIDELINES

Executive Summary

With the advent of the technological revolution and the digital transformation that made all scientific disciplines becoming computational nowadays, High-Performance Computing (HPC) is increasingly identified as a strategic asset and enabler to accelerate the research performed in all areas requiring intensive computing and large-scale Big Data analytic capabilities.
The University of Luxembourg (UL) operates since 2007 a large academic HPC facility that remains the reference implementation within the country, offering a cutting-edge research infrastructure to Luxembourg public research while serving as edge access to the upcoming Euro-HPC Luxembourg supercomputer. Special focus was laid on the development of large computing power combined with huge data storage capacity to accelerate the research performed in intensive computing and large-scale data analytic (Big Data).
For more details: see hpc.uni.lu (Main contacts: Prof. Pascal Bouvry (Head), Dr. Sébastien Varrette (Deputy head), HPC for research).

The University extends access to its HPC resources (i.e., facility and expert HPC consultants) to the scientific staff of national public organizations and external partners for the duration of joint research projects under the conditions defined in this document.

Version 1.0 - Approved by UL rectorate on July 7th, 2020

Overview of the UL HPC Facility.

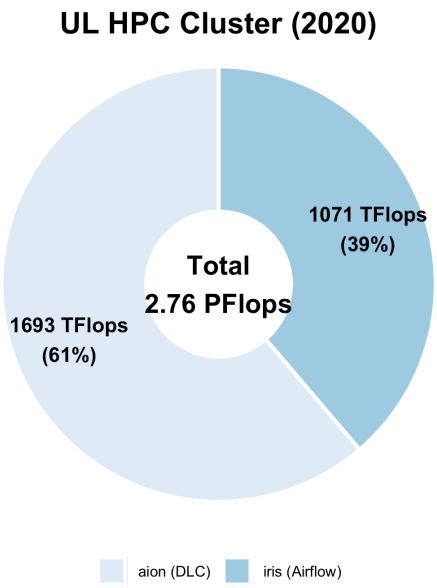

The UL HPC platform has kept growing over time thanks to the continuous efforts of the core HPC team (Dr. S. Varrette, S. Peter, H. Cartiaux, Dr. F. Pinel, Dr. E. Kieffer, Dr. E. Krishnasamy, T. Valette, A. Olloh - contact: hpc-team@uni.lu).
Installed in the premises of the University's Centre de Calcul (CDC), it provides in 2020 a total computing capacity of 2.76 PetaFlops (1 PetaFlops = 10^{15} floating-point operations per second) across two clusters of compute nodes called *iris* and *aion*, and around 10 PetaByte of shared data storage. This places the HPC center of the University of Luxembourg as one of the major actors in HPC and Big Data for the Greater Region Saar-Lor-Lux. It also consolidates the University's ambition to offer a cutting-edge research infrastructure to Luxembourg public research while serving as edge access to the upcoming Luxembourg MeluXina supercomputer in the EuroHPC context.

In practice, the UL HPC Facility features 3 types of computing resources:

- "*regular*" nodes: Dual CPU, no accelerators, 128 to 256 GB of RAM
- "gpu" nodes: Dual CPU, 4 Nvidia accelerators, 768 GB RAM
- "*bigmem*" nodes: Quad-CPU, no accelerators, 3072 GB RAM

These resources can be reserved and allocated for the execution of *jobs* scheduled on the platform thanks to a Resource and Job Management Systems (RJMS) - Slurm in practice. This tool allows for a fine-grain analysis and accounting of the used resources, facilitating the generation of activity reports for a given time period.

Jobs are scheduled on computing resources depending on their purpose explicitly specified upon submission i.e., interactive tests, normal 'batch' run on the above-mentioned type of resources, long term executions on regular nodes. A job is characterized (and thus billed) according to the following elements:

- `NNodes:` number of computing nodes
  and, per node:
  - `Ncores`: number of CPU cores allocated *per node*
  - `Mem`: memory size allocated *per node,* in GB
  - `Ngpus`: number of GPU allocated *per node*
- execution time (in hours) (T_{exec})

Main Terms and conditions of HPC Resource Allocations for Research Projects and External Partners

Detailed conditions will be defined in the legal contract that will have to be signed by the project PI, the funding agency and the UL upon project acceptance. It is adapted from the existing HPC service agreements already established with external industrial partners allowing for accessing the UL HPC resources i.e., facility and expert consultants. The proposed cost model is the same for all external partners. With regards to private partners, no more than 10% of the available computing capacity can be allocated for their usage.
Here are the main budget elements

1. HPC infrastructure access is granted for the time of the project for the set of Personnel contributing to the project (Project End Users).
    a. All Project End Users will have to conform to the University HPC user charter as defined on https://hpc.uni.lu/users/AUP.html;
    b. For usage accounting purposes, their University HPC accounts will be hierarchically linked to (1) the Project name (2) the project PI and, when applicable, to the employing organization.
2. The utilization of the University computational resources is charged according to the following conditions:
    a. One Service Unit (SU) equals 1 hour on 1 physical processor core (with 4 GigaByte RAM per core available) on a *regular* computing node with no accelerators or large memory capacity) and is charged at the price of 0,03€ per SU (VAT excluded).
    b. To determine the exact number of SU associated with a job, additional weighted factors $\alpha_{<type>}$ are taken into account in an automatic way by the scheduler:
        - α_{cpu}: the normalized relative performance of the CPU processor core (reference: peak performance of skylake processor on regular computing nodes i.e. 73,6 GFlops/core)
        - α_{mem}: inverse of the average available memory size per core
        - α_{GPU}: weight per GPU accelerator

        Then the number of SU associated to a job is defined with the following formulae:

        $$\mathtt{NNodes} \times [\alpha_{cpu} \times \mathtt{Ncores} + \alpha_{mem} \times \mathtt{Mem} + \alpha_{GPU} * \mathtt{Ngpus}] \times T_{exec}$$

        The values of the factors depend on the resources used. The below table reviews the implemented values expected for the time period 2021-2022:

Cluster	Type of node	Partition / Usage	#Cores per node	CPU	α_{cpu}	α_{mem}	α_{GPU}
Iris, Aion	regular	interactive	28/128	n/a	0	0	0
Iris	regular	batch, long	28	broadwell	1.0*	1/4 = 0,25	0
Iris	regular	batch,long	28	skylake	1.0	1/4 =0,25	0
Iris	gpu	GPU	28	skylake	1.0	1/27 ≈ 0,037	50,0
Iris	bigmem	Large-memory	112	skylake	1.0	1/27 ≈ 0,037	0
Aion	regular	batch, long	128	epyc	0,57	1/1.75 ≈ 0,57	0

** the performance ratio cannot be adapted in this case for technical reasons.*

Example of estimated usage for budget planning:

- Continuous use of ***2 regular skylake nodes (56 cores, 224GB Memory)*** on iris cluster (28 cores per node, 4 GigaByte RAM per core i.e., 112GB per node) during a period of ***30 days***:
  2 nodes * [(α_{cpu} + α_{mem} x 4) x 28 cores] * 30days * 24hours =
  2 nodes * [(1.0 +1/4*4)*28 cores] * 720 hours =
  80640 SU = 2419,2€ VAT excluded
  for a period of ***60 days***: 4838,4€ VAT excluded

- Continuous use of ***2 regular epyc nodes (256 cores, 448GB Memory)*** on aion cluster (128 cores per node, 1,75 GigaByte RAM per core i.e., 224 GB per node) during a period of ***30 days***:
  2 nodes * [(α_{cpu} + α_{mem} x 1,75) x 128 cores] * 30days * 24hours =
  2 nodes * [(0.57 +1/1.75*1.75)*128 cores] * 720 hours =
  289382,4 SU = 8681,47€ VAT excluded
  for a period of ***60 days***:17362,94€ VAT excluded

- Continuous use of ***1 GPU nodes (28 cores, 4 GPUs, 756GB Memory)*** on iris cluster (28 cores per node, ***4 GPUs per nodes***, 27 GigaByte RAM per core, 756 GB per node) during a period of ***60 days***:
  1 node * [(α_{cpu} + α_{mem}*27) x 28 cores+ α_{GPU}*4 GPUs]* 60days * 24h=
  1 node * [(1.0 + 1/27*27) x 28 cores+50.0*4 GPUs]* 1440 hours =
  368640 SU = 11059,2€ VAT excluded
- Continuous use of ***1 Large-Memory nodes (112 cores, 3024GB***

Memory) on iris cluster (112 cores per node, 27 GigaByte RAM per core i.e. ***3024 GB per node***) during a period of ***60 days***:

1 nodes * [(α_{cpu} + α_{mem} x 27) x 112 cores] * 60days * 24hours=
1 nodes * [(1.0 +1/27*27)*112 cores] * 1440 hours =
322560 SU = 9676,8€ VAT excluded

If the project PI is not able to anticipate the type and amount of resources needed, **we suggest to account 5529,60€ for every 12 PM of funded personnel** (i.e., 1 month of continuous usage on the most expensive type of resource).
In particular, by default:

- ❏ Budget for 1 funded PhD student (36+12PM): 22118,4€
- ❏ Budget for 1 funded PostDoc (24PM): 11059,2€

3. Computing capacity allocation is done on a fair-share principle, with no guarantee of being satisfied. Note that general statistics on the platform demonstrate in the past year (i.e., 2019 on iris) a yearly availability of 96.69% for an average load usage of 65,2% and a median slowdown of 1.0038 in 2018. Due to the relatively low number of specialized nodes (i.e., GPU and Large-Memory), their access is restricted and prioritized for UL internal projects.
4. An estimation of the Project's expected resource utilization during its duration is to be provided at the start of the contract
5. A data storage project directory (with a capacity of 1 TeraByte) is created for the Project free of charge.
    a. Capacity **extensions are possible and will be charged at the price of: 100€ (VAT excluded) per each additional TeraByte per month**
    Example: 5 additional TeraBytes (for a total of 6TB available) for 36 months:
    5TB * 36months * 100€ = 18000€ VAT excluded.
    b. All data stored on the UL HPC facilities by the Project End Users remain under the responsibility of the project PI. As per GDPR rules, the Project PI is informed that at least one copy of the data, for disaster recovery purposes only, will be stored in University facilities. The Project PI can request the deletion of this data at any moment.
6. Expert-level service and support can be provided by the University HPC staff in the form of pools of 4-hour specialized help, at 480€ VAT excluded (equivalent to 120€ per hour expert rate EA). Service and support activities

include but are not limited to HPC training on facilities utilization, software installation, HPC workflow development. If HPC expertise consultants are required on a project, they need to be included under the 'Staff' heading in FNR projects at the submission deadline

7. Usage of the platform for the Project accounting will be reported on a yearly basis.
    a. The report will include the number of CoreHours and Service Units utilized (aggregate and per Project user).
    b. The report will include the storage space utilization
    c. the appropriate VAT rate will be added where applicable
8. Billing will be performed on a yearly interval and shall be paid by the Project PI upon presentation of a proper invoice issued by the University, following the instructions stipulated on this invoice (bank details, communication on bank transfer...).
9. For auditing purposes, a final report summarizing the utilization of the HPC Resource allocation over the project lifetime will be issued.