

The ISC Computational Center

We are currently **test driving** a strategy for mutualizing computational resources within HEI. The current policy is therefore designed to be temporary and will evolve as we understand better the usage patterns and requirements of various users.

The project is expected to have three phases, which are subject to change:

- **Initial phase** (until March 2025): we run a best-effort, priority-less offering. The goal of this phase is to test-drive some of the tools we use (Slurm, Apptainer, etc.), estimate the demand and requirements from users, and collect data on usage patterns, e.g., periods of congestion and overall usage ratio.
- **Expansion phase** (April-July 2025): we introduce priorities based on investments and projects' criticality and we stabilize the toolchain, with an improved documentation and more formal support to the users. We formalize and possibly extend the scope of the offering to more research labs and institutes within HEI.
- **Exploitation** (August 2025-onwards): depending on the success of earlier phases, we introduce service-level agreements (SLA) and a tarification for users to contribute financially to the platform based on their usage patterns.

If you need access, please use the [following form](#) to make a request. Once your request has been received, we will either provide you with access immediately or organize a meeting to discuss any specific requirements you might have.

The following rules apply in this initial phase of the project:

- All users must submit [Apptainer jobs](#) via Slurm. The rationale for this restriction is to ensure we keep visibility on the cluster usage and prevent maxing out cluster resources. **A proper documentation will be made available soon**
- Users must regularly clean the Apptainer cache using `apptainer cache clean`.
- Users are expected to search for solutions and contribute to the documentation whenever possible. Remember that this is a best-effort service without any guarantee of availability and/or support.
- There is a limited prioritization: in the future, users will be assigned tokens depending on criteria to be approved by the steering committee. During the initial phase, a manual tuning will be performed. Please note that we don't really expect this constraint to be a problem, as the current number of users is limited.

Computational Resources

The available computational resources currently available are composed of

- [The Calypso cluster](#)
- [The Rumba production servers](#)
- [DISCO and ChaCha supercomputers](#)

Calypso

It is currently composed of the following machines :

- 1 DELL R740XD (Master)
- 12 DELL R630 (6 currently active)
- 3 DELL R630 (1 spare in 23N321, storage 1)
- 1 DELL 7920 (in Rumba, currently shutdown)

It can be accessed via a [Wireguard VPN](#) (ask Rémi for access).

Rumba

Those are the production servers, running various ISC services (such as Moodle, this Wiki and other tools). Feel free to ask for more if required!

Disco & ChaCha

Research servers, equipped with 2x NVIDIA A100 (DISCO) and 2x NVIDIA H100 (96 GB of RAM each) (ChaCha), respectively. Documentation for accessing those machines is available here [Using the research servers DISCO and ChaCha](#)

VPS

We currently have 2 Virtual Private Server (VPS) hosted at [Infomaniak](#).

1. Hannibal is the VPS hosting the [ISC Learn Moodle platform](#)
2. Marcellus is the legacy VPS hosting various services, such as
 - XXX
 - YYY

From:

<https://wiki.isc-vs.ch/> - **The ISC wiki**

Permanent link:

<https://wiki.isc-vs.ch/doku.php?id=administratif:compute&rev=1733478954>

Last update: **2024/12/06 09:55**

