



**Nyx, a new high
performance
server at the ICC**

*In Greek mythology, Nyx is the goddess
and personification of the night.*

Goal: we want to establish a strong foundation for a computer cluster which can be used by several groups at the ICCUB. Sharing this resource will achieve several goals:

- 1. Better efficiency of use:** the computing needs of a single group or PI fluctuate in time, peaking at the “research and production of results” phase but decreasing dramatically when in “paper writing” mode. Different groups will naturally have this research cycle staggered thus ensuring a more constant and efficient use of the infrastructure;
- 2. Maintenance and future extensions** is shared among different groups, which also makes it more efficient;
- 3. It makes ICCUB a more attractive host institution** for fellowships, international visitors and ICREA staff and it opens doors to new funding streams because it enables research that is out of reach with the obsolete machine;
- 4. Sharing infrastructure across different research groups fosters cross-field collaborations and encourages cross-disciplinary synergies.** Building bridges across different research directions is extremely important to facilitate innovation; although it is difficult to predict the outcome of such collaboration, it is undeniable that cross-fertilization of ideas across fields is one of the main drivers of innovation.


Initial participating groups

	Share	IP
CosmoGal	42%	C. Laporte
PhysCosm	16%	L. Verde
ICC	14%	ICC
GravWaves	16%	M. Gieles
LIA-8	12%	X. Luri

Two more groups are expected to join shortly. More groups can join in the future (up to the physical limits of the setup).

New groups are expected to contribute to the common infrastructure/resources.

The server



11 nodes

1408 cores

22TB RAM – infiniband shared

352TB qNAP disk system (max 1PB)

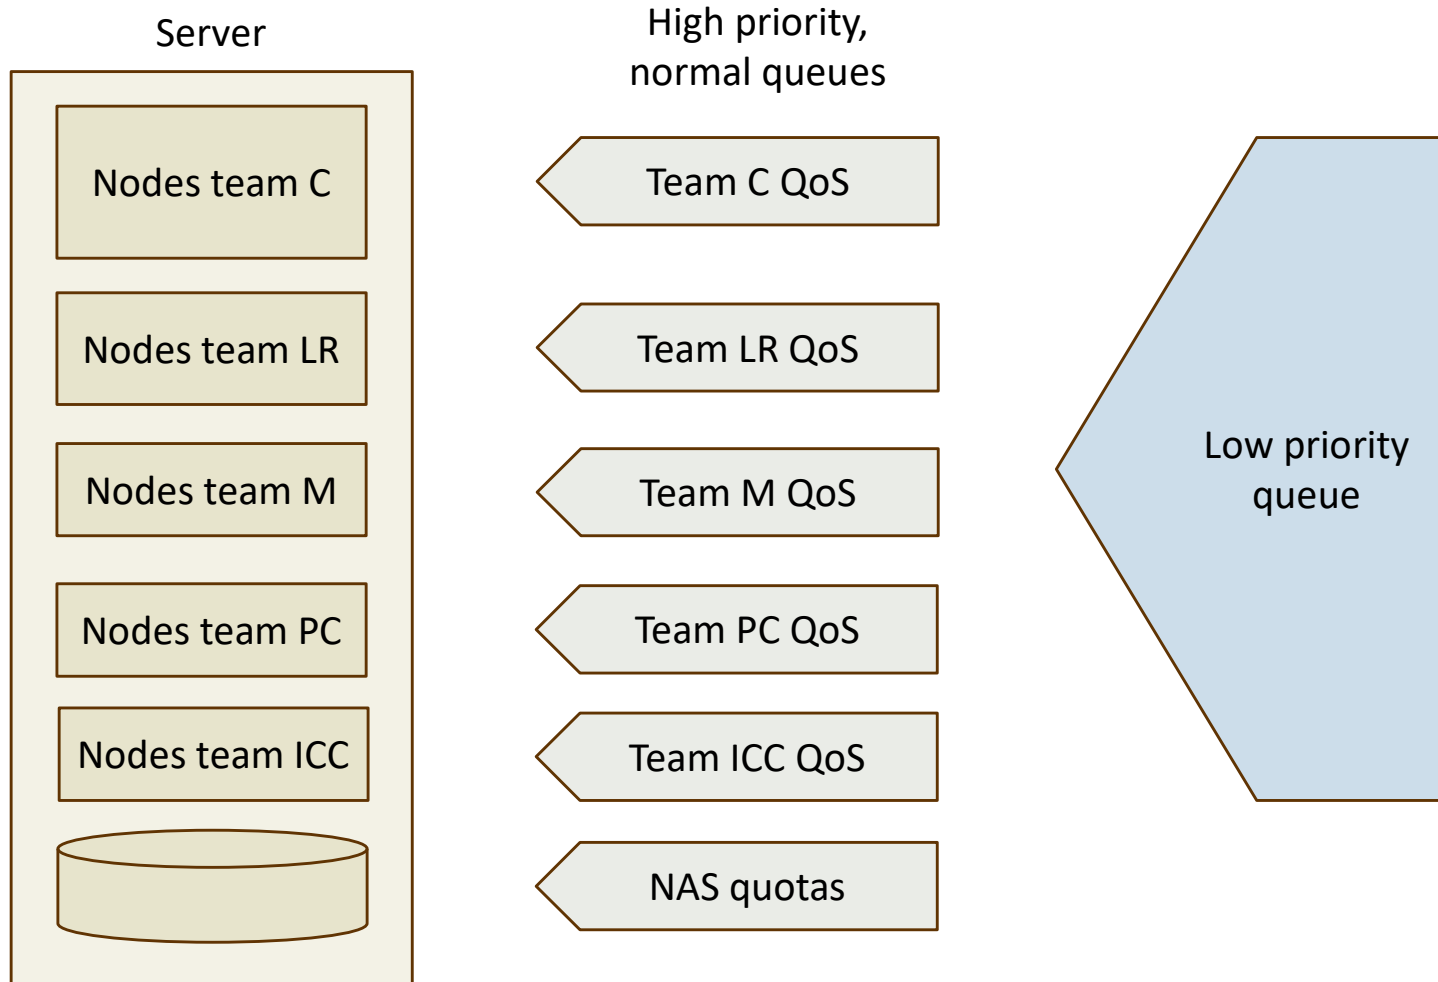
Located at the new computer room of the faculty of physics

Configuration based on QoS/queue system

The server is not physically segmented:

- Exclusivity of use of resources can be rearranged at any time by queue/QoS configuration and disk quotas. If after some months we are not happy with the results, we can change the configuration!
- Allows for the use of any fraction of the server, up to the full set of nodes
- Simplifies operations, allows a more efficient use of resources.

Initial resource allocation concept



Teams have priority for the use of their share of the machine

- The team can only use N_{team} cores + M_{team} memory simultaneously.
- Access to team resources with top priority through the normal queues (see later). In normal circumstances should provide a fair access to the team resources.
- Preemption priority also possible if needed: immediately stop or halts any other job running to get the requested cores. This is applied in a limited form in the low priority queue; wider use would be disruptive and not recommended.
- Alternative approach preferred: with a forewarning of need of exclusivity, cores can be taken out of the general pool for a given period, leaving them only available to the corresponding team for the requested duration.
- Specific access to GPU nodes through slurm “Features”, with enhanced priority to ensure easy access to these nodes.

ICC team access

Special definition of QoS for ICC team members:

1. **Test QoS:** limit of 16 cores. Used for development and learning.
2. **Medium QoS:** limit of 48 cores for normal jobs.
3. **Normal team QoS:** limit of 96 cores (full ICC share) for authorised jobs.

Access to ICC resources to be organized by the informatics commission FQA/ICC

Access beyond team boundaries

- A lower priority queue (how low TBD). Allows access beyond team limits at a lower priority and with preemption; thus, resources from other teams can be used with less priority than its owners.
- Two types of users:
 - Normal users (teams)
 - **Invited users (ICC)**

The access of invited ICC users to Nyx will be defined under the following guidelines:

- The spare time per period (e.g. semester) will be estimated based on the actual use of the machine
- ICC members will be invited to apply for the use of this spare time; a formal project documenting the need for High Performance Computing and the suitability of Nyx will be required
- A committee formed by representatives of the funding groups and the ICC will allocate the spare time among the invited users each period
- This allocation has to be understood to be done in a “best effort” basis, without guarantees of execution in a given period.

Nyx is now being set up. Expected to be fully operational in about a month.