

Research Computing Resources Allocation

Purpose

This Policy aims at providing a guide for fair allocation of research computing resources in XXXX that will be used by the XXXX Research Computing Allocation Committee (RCAC). The resource allocation policy is defined with the objective of fairness and efficiency in use to achieve XXXX's research and academic goals. The policy defines resources for allocation, queues, disk quotas and backups, archiving, outside network connectivity, physical space allocation in computer rooms and governance that could affect research computing.

DEFINITIONS

Shared resources

Shared research computing resources covered by this policy fall into two categories:

- Shared: XXXX resources intended for general shared use (e.g. compute systems, infrastructure, etc.).

- Semi---shared: Resources that are specifically allocated to entities within XXXX, but are intermittently (e.g., when not in use by the principal group(s), etc.) made generally available for research computing as governed by an agreement.

Users

Users in research computing we will categorize as:

- **Group users: research centers, project groups** (may include external

collaborators), etc. These are led by PIs who are responsible for research computing allocations within groups' quotas.

- Faculty.

- Students (for their normal work towards completion of degrees).

- Production units, such as core labs.

- External users without XXXX collaborating partners.

ALLOCATIONS

□ GPU nodes: allocable to approved users only.

The minimum allocable unit is one compute core within a node, although entire nodes can be specifically requested in justified cases.

Queues

In order to maximize the simplicity of resource allocation in research computing, the allocation tasks will be initially delegated to an appropriate scheduling system that implements comprehensive scheduling policy. This approach removes the need to define additional queues, though they may be warranted in special circumstances only by approval from RCAC.

For servicing jobs on specialized hardware, such as high memory nodes, a combination of resource requirements and pre-emptive job scheduling will be used to control this type of jobs. This idea applies to all nodes that have special resources (GPU, SMP, high memory). The scheduler will pre-empt (terminate) jobs from non-approved users should an approved user submit a job that needs the resource. Non-approved users may specify if pre-emption of a job is acceptable in the associated batch script. If not, the job will not run on nodes that have special resources. The performance indicators of all allocations will be carefully monitored and as the result the queuing policy may be modified. The RCAC committee will decide on the need to change the queuing policy based on monthly utilization reports.

Debugging queue ? Comments

Interactive queue ? Comments

Resource allocation policy

XXXX implements as a part of fair use policy a **user quota accounting system** for allocation of research computing resources. The accounting is based on the concept of allocable service unit (ASU) representing an allocation of a single core for one wall-time hour (adjusted for relative processing powers when dealing with non-homogenous environments).

In order to streamline the policy, we envisage users submitting jobs via two modes:


Resource Managing Units (CRMUs) that manage the resources, depending on the size of the request.

For example, if the total capacity of the system is C cores, resulting in $T = C * 365 * 24$ ASU per year, then 10% of these ASU ($0.1 * T$ ASU) will be made available for automatic allocation. Allocations per student per year and per PI per year will be determined based on the projected number of students and PI that will be using the resources for the year in question.

Once the allocated credits to a user are exhausted, subsequent jobs submitted by

such users will still be allowed to execute, but strictly on the best effort basis, with fixed default (lower) priority. Alternatively, such users can request additional quotas from RCAC with proper justification.

The principle of group allocation is shown in the diagram below. Groups are free to determine individual allocations of group members so as group's usage does not exceed the aggregate value allocated to the group. Users must specify to which allocation group jobs are being submitted so that utilization of resources can be properly debited from the correct group's quota.



utilization of these is counted against the users' quotas.

External users without XXXX collaborating partners can utilize the research

computing resources only based on the written agreement with XXXX approved by

RCAC and if such use will not affect academic users in any significant manner.

Disk storage and backup

Default allocation on shared systems: