

SLURM configuration

This section details the current SLURM configuration for the ISC Computational Center.

Installation

SLURM has been installed from tarball, version : 24.11.0

All official plugins installed :

- libnvidia-ml
- TODO

Install / upgrade process TODO

Architecture

Chacha

- Client (slurm-smd-client)
- Worker (slurm-smd,slurm-smd-slurmd)
- Controller (slurm-smd-slurmctld)
- Accounting DB (slurm-smd-slurmdbd)

Disco

- Client (slurm-smd-client)
- Worker (slurm-smd,slurm-smd-slurmd)

Schema

TODO

Partitions

Dance

This is the default partition : is is currently composed of Chacha and Disco.

Chacha and Disco

These partitions can be used to restrain an account to use only one server.

More details of partitions usage TODO

Accounts

Accounts have been created in 2 groups :

- Premium Researchers (premium_rs) :
- Standard Researchers (standard_rs) : All users who can't participate financially to the project. Students are also part of this group

There are 2 other groups : Test and temp : Test is only for administration purpose, and temp is a locked group either to migrate someone from another account (can't delete an account when someone has it as a default account) or to disallow someone to run jobs (MaxSubmitJob=0)

TODO

QOS and Limits

Current limits on QOS and accounts applied to each project account :

- **premium_rs** :
 - MaxCPUs=44
 - MaxNodes=2
 - MaxTRES=gres/gpu=1,gres/shard=96,cpu=44,mem=500G
 - GrpTRES=gres/gpu=1,gres/shard=96,cpu=44,mem=500G
 - GrpWall=3-00:00:00
 - MaxWall=3-00:00:00
- **standard_rs** :
 - MaxCPUs=24
 - MaxNodes=1
 - MaxTRES=gres/gpu=1,gres/shard=96,cpu=24,mem=256G
 - GrpTRES=gres/gpu=1,gres/shard=96,cpu=24,mem=256G
 - GrpWall=1-00:00:00
 - MaxWall=1-00:00:00

TODO : complete / modify until limits are finished

Scheduling

Fairshare is one way of prioritizing jobs in the job queue.

- premium_rs :
 - Fairshare : 750
- standard_rs :
 - Fairshare : 250

TODO

User creation

Users fill their form from the page [ISC Computational Center](#) , then their information is used to create their SSH access, and their SLURM user in its project account, according to the SLA / QOS we can provide : Premium or Standard.

Several users can share the same project account to work as a team. (Limits are applied both to each user for maximum limits and the group for group limits)

Currently, Rémi creates and configure user accounts.

User creation process TODO

Backups

Currently, all the SLURM configuration is manually backuped (files, DB)

TODO : automate and redirect to backup server when its ready

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

Permanent link:
<https://wiki.isc-vs.ch/doku.php?id=infra:compute:slurmconfig&rev=1746790698>

Last update: **2025/05/09 11:38**

